Web Semantics: Science, Services and Agents on the World Wide Web 35 (2015) 25–34

Contents lists available at ScienceDirect

Web Semantics: Science, Services and Agents on the World Wide Web

journal homepage: www.elsevier.com/locate/websem

Querying and integrating spatial–temporal information on the Web of Data via time geography

Carsten Keßler [a,*], Carson J.Q. Farmer [b]

[a] Center for Advanced Research of Spatial Information and Department of Geography, Hunter College, City University of New York, 695 Park Avenue, New York, NY-10065, USA

[b] Department Geography, University of Colorado at Boulder, Boulder, CO-80309, USA

ARTICLE INFO

Article history:
Received 16 February 2015
Received in revised form 28 September 2015
Accepted 28 September 2015
Available online 3 November 2015

Keywords:
Time geography
Web of Data
Ontology design patterns
Data integration
Linked open data
Federated querying

ABSTRACT

The *Web of Data* is a rapidly growing collection of datasets from a wide range of domains, many of which have spatial–temporal aspects. Hägerstrand's *time geography* has proven useful for thinking about and understanding the movements and spatial–temporal constraints of humans. In this paper, we explore time geography as a means of querying and integrating multiple spatial–temporal data sources. We formalize the concept of the space–time prism as an ontology design pattern to use as a framework for understanding and representing constraints and interactions between entities in space and time. We build on a formalization of space–time prisms and apply it in the context of the Web of Data, making it usable across multiple domains and topics. We demonstrate the utility of this approach through two use cases from the domains of environmental monitoring and cultural heritage, showing how space–time prisms enable spatial–temporal and semantic reasoning directly on distributed data sources.

1. Introduction

The World Wide Web is currently undergoing a rapid change from an information resource primarily targeted at human users towards a distributed knowledge base that provides structured, machine-readable information. This *Web of Data* covers an increasing breadth of topics, ranging from museum collections [1] and drug databases [2] to humanitarian aid data [3] and real-time sensor data [4]. Many of these data sources follow the Linked Open Data (LOD) principles [5] that allow for the integration of initially isolated *data silos*. The number of applications leveraging this large distributed database is constantly growing, including examples as diverse as the BBC program Web pages [6], drug discovery tools [7], university management and campus navigation apps [8,9], and exploration and visualization tools for rainforest preservation data [10].

A large number of resources on the Web of Data have spatial or even spatial–temporal components. This information can be recorded at a very coarse level, such as the birth and death dates and places of a person[1] or the production year and place of a cultural heritage artifact.[2] At the other end of the spectrum, such data can be of high granularity, such as photo metadata[3] or measurements provided via Sensor Web services [4]. In these cases, location is often recorded at sub-meter resolution and timestamps go down to the minute or even second. Along this spectrum, these spatial–temporal characteristics have the potential to identify interactions between objects that co-occur in space and time.

The conceptual framework of time geography [11] has been introduced to model the constraints that restrict humans in moving and interacting with each other (and their environment). Building on the elegant and intuitive notions of space–time paths and prisms, time geography has been applied to a number of different use cases that all involve capturing and planning joint *human* activities [12–17]. In this paper, we broaden this view to analyze potential interactions between any spatial–temporal objects. The goal is to utilize time geography as a framework for data integration based

* Corresponding author.
E-mail addresses: carsten.kessler@hunter.cuny.edu (C. Keßler), carson.farmer@colorado.edu (C.J.Q. Farmer).

1 http://d-nb.info/gnd/11855042X, for example.

2 http://www.kulttuurisampo.fi/item.shtml?itemUri=http%3A%2F%2Fkulttuurisampo.fi%2Fannotaatio%23Instance_esine_H84039-21156597532081, for example.

3 See http://500px.com/photo/61774497 as an example; while technically not Linked Data, the page contains the photo metadata in a machine-readable format.

http://dx.doi.org/10.1016/j.websem.2015.09.005
1570-8268/

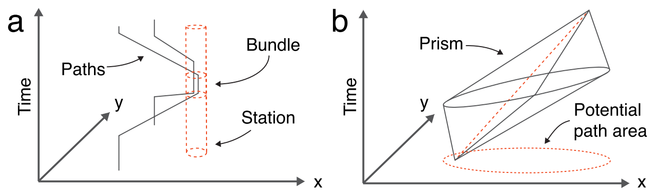

Fig. 1. Features of Hägerstrand's time geography. (a) Space–time paths with bundling at a station; (b) Space–time prism with accompanying potential path area. *Source:* Adapted from [23,24].

on the spatial and temporal characteristics of information objects. In this context, *data integration* refers both to the reconciliation of different identifiers for the same real-world objects, as well as the identification of *related* objects. In environmental monitoring, for example, objects related to the measurement of a given pollutant could be other measurements of the same or related pollutants in the vicinity, or nearby events that may have influenced the measured values, such as a traffic jam, a fire, or an extreme weather event.

To enable the systematic identification of such related objects, we introduce an ontology design pattern [18] that formalizes the space–time prism using the Web Ontology Language (OWL). It builds on previous work on semantic trajectories [19] and integrates it with the measurement theory for time geography introduced by Miller [20]. The formalization presented in this paper allows for the application of time geography concepts directly on the Web of Data.

The remainder of this paper is organized as follows: The next section reviews relevant related work, followed by a generalization of time geography in Section 3. Here we outline how space–time prisms may be used to integrate disparate data sources. Section 4 discusses an ontology design pattern for space–time prisms, followed by application examples in Section 5. Section 6 presents conclusions and a brief discussion of future work.

2. Related work

This section points to related work on time geography and ontology design patterns.

2.1. Time geography

Hägerstrand's conceptual framework for *time geography* [11] was originally developed to understand how human migration activities are constrained at the individual level. Recently, there has been renewed interest in time geography concepts in research using geographic and/or spatial data (see [21,22] and references therein). The basic concepts are the *space–time path* (see Fig. 1(a)), showing an individual's location changing with time, and the *space–time prism* (see Fig. 1(b)), showing an individual's travel potential. This potential is constrained by the speed at which the individual can travel, as well as locations at which he or she has to be present, such as home or work.

Time geography has been applied in a wide range of scenarios, including location-based services for individuals [12,13] and user groups [14], accessibility models [25], shared-ride trip planning [15,16], delineation of wildlife home ranges [26], gender differences in access to opportunities [27], understanding place in contextual models of health [28], and geovisualization of human activity patterns [29]. Raubal [30] uses time geography to track concepts – instead of people – through time, combining the framework with conceptual spaces [17]. Time geography has since been formalized in different ways, most thoroughly by Miller [20], who introduces a rigorous measurement theory that formally defines the main elements based on the basic assumptions that: (1) the metric space satisfies the notions of identity, non-negativity, and triangular inequality about distance; (2) data are recorded at specific points in time; and (3) we have perfect information about the system (although Miller [20] and others [31,32,24] have explored relaxation of this assumption to some degree). Building on these relatively simple assumptions, Miller [20] develops a measurement theory for time geography, with mathematical definitions for space–time paths, prisms, stations, bundles, and intersections.

A space–time *path* consists of a sequence of control points and a corresponding sequence of path segments connecting these points. In this definition, control points are observed or measured locations in space and time, and segments connect temporally adjacent control points (usually represented as a straight line between observed points). A space–time *prism* may exist between any pair of temporally adjacent control points, creating an open temporal interval during which un-recorded travel may occur. An object may thus occupy locations in space other than the straight-line segment between two adjacent control points (see Fig. 1(b)). Stations are generally designated activity locations (e.g., home or office), and are sites where multiple paths can bundle (convergence of two or more paths for some shared activity) for some given length of time. This has direct relevance to space–time intersections, which is the case of two or more time geographic features sharing some locations in space and time. Space–time intersections are generally more explicit than bundling and may be the result of a shared space–time activity.

Miller [20] provides strict conditions within which space–time paths are bundled and where intersections may occur between paths and prisms; he also provides geometrically-based definitions for paths, prisms, and related constructs. While we also adopt this traditional object-based view on time geography, these strict conditions can be *relaxed* in some sense by allowing for imperfect information or non-uniform velocities [33]. By placing time geography within a stochastic framework, qualitative statements about *potential* interactions within intersecting prisms can be quantified using *a priori* knowledge about an object's behavior and some knowledge of the probability distribution of an object's location over time [31,32].

2.2. Ontology design patterns

The concept of a *design pattern* has its origins in architecture and has been adopted for software engineering, where it refers to a template solution for a recurring software design problem [34]. Such patterns can address different aspects of software design, such as structural design, algorithm strategies, or user interaction. Ontology design patterns (ODPs) [18,35] transfer this idea from software to ontology engineering. They introduce template solutions for common ontology modeling problems. An instance (i.e., an actual implementation) of an ODP can typically be found as a component or module of an ontology.

ODPs can be classified into *logical* and *conceptual* patterns [36]. While the former focus on the formal specification of a pattern at the logical level (e.g., in first order or description logic), the latter generally target the conceptualization of an ontology fragment. The participation pattern extracted from DOLCE [37] is an example of such a conceptual ODP. It describes the relationship between an object located in space and the event it participates in, which is temporally located in a time interval (see Fig. 2).

Ontology design patterns that explicitly address the handling of spatial information include patterns for spatial data quality [38], for quantification over types (e.g., to count the number of species in a region) [39], for cartographic map scaling [40], and for the relationships between stimuli, sensors, and observations [41]. In the context of time geography, recent research on the semantic

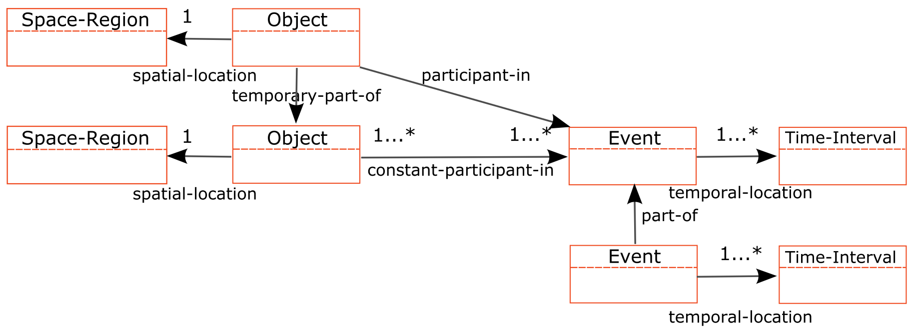

Fig. 2. Schematic overview of the participation ontology design pattern.
Source: Based on [36].

enrichment of trajectories [42,43] demonstrates how to attach semantic data to specific parts of space–time paths. This can greatly improve discovery, reuse, and integration of trajectory data from different sources. Hu et al. [19] have built upon the notion of annotated trajectories by defining a formal ontology design pattern for semantic trajectories that is directly applicable to a variety of trajectory datasets and easily extensible to other domains or application areas. In Section 4, we build upon this work from the perspective of time geography in order to develop our own space–time prism design pattern for spatial–temporal data integration.

3. Generalizing the space–time prism

The measurement theory presented by Miller provides the necessary tools to develop a framework for data integration across a number of fields and data sources. Indeed, Miller alludes to the idea of time geography as a means to develop an efficient query design language in [20]. In the following, we propose a novel perspective on the space–time prism that allows us to apply selected time geographic principles for data integration and query design.

3.1. Spatial autocorrelation

"Everything is related to everything else, but near things are more related than distant things" [44]. This idea is well-known as Tobler's First Law of Geography and also referred to as (positive) spatial autocorrelation. In the following, we will consider the relation between spatial autocorrelation and time geography and the space–time prism in particular. The interactions we consider here result in *predefined* and *inferred* space–time prisms. We will discuss these two kinds of prisms in Section 3.2 as a motivation for the formalization presented in Section 4.

Spatial autocorrelation essentially quantifies the correlation of a variable with itself through space. This being the case, systematic patterns in the spatial distribution of a given variable are often the result of spatial autocorrelation, be it positive (i.e., nearby regions display similar values), negative (i.e., nearby regions display dissimilar values), or zero (i.e., no significant association between regions) spatial autocorrelation. In this particular analysis, we utilize the Moran's I measure of spatial autocorrelation, which is designed to compare spatially *local* averages with *global* averages of a given variable by considering both the locational and attribute relationship between each control point i across all space–time objects and its surrounding neighbors j [45].

The metric is similar to the usual correlation coefficient, with the addition of a spatial weights matrix with elements W_{ij}, such that entries in the matrix are coded as one if two control points are 'neighbors' and zero otherwise. Here, neighbors are defined as all control points within a given distance d of control point i, or alternatively, some given number of nearby neighbors k (alternative definitions of the weights matrix [and neighbors] are possible). In order to derive a measure of spatial autocorrelation at every control point across all trajectories, a *local* measure of spatial autocorrelation must be used (see [46] for a comprehensive discussion of these metrics).

3.2. Interpretations of the space–time prism for data integration

In classical time geography, the shape of a space–time prism is defined by the maximum travel velocity and the space–time stations at the beginning and end of the time period available. If we generalize the maximum travel velocity to any kind of interaction in space, any measure of 'relatedness' can be used to define the shape of the prism for a data integration task; this includes using the degree of spatial autocorrelation of the phenomenon under consideration.

In the simplest case where some measure of 'relatedness' is known, for any given recorded spatial–temporal location of the object under consideration (i.e., a *control point* measured in space and time), we can identify all potentially related objects by checking whether they can be found within some *predefined* space–time prism. Fig. 3(a) illustrates this case. In air quality monitoring, for example, the size of the prism can be defined by the dispersion rate of the pollutant under consideration. Objects inside the prism could be sources for the pollutant that are reflected in our measurement.

The phenomenon we are looking at may not always be as well understood as in the dispersion case, or we may be looking at a set of spatial–temporal objects that we know are related to one another, but we do not know to what extent (or how). In this case, the space–time prism can be *inferred* from the data, and the degree of spatial autocorrelation between the objects – i.e., the size of the prism – can tell us something about how *local* the phenomenon at hand is. Consider the example of recording exhaust fumes produced by cars moving through a city (see Section 5 for details): A cluster of peaks in fumes exhausted by different cars may point to a local reason such as a speed bump, whereas non-clustered data would attribute the peaks to non-spatial reasons such as different driving styles or vehicle types.

4. An ontology design pattern for the space–time prism

In the following, we introduce an ontology design pattern for the space–time prism as one of the central concepts in time geography, based on the discussion in Section 3. We also demonstrate the pattern's compatibility with different existing ontologies and design patterns that define different kinds of interactions between the spatial and temporal dimensions.

4.1. Motivation

Location in space–time allows us to ask many interesting questions concerning the nearness and relatedness of objects in space, as discussed in the previous section. In order to utilize

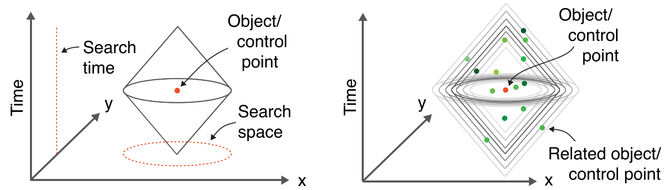

(a) Predefined prism. (b) Inferred prism.

Fig. 3. Predefined (a) and inferred (b) space–time prism for data integration.

this potential, a formalization is required that is both rigorous enough to support the required mathematical operations, yet flexible enough to work on a broad range of existing datasets. Ideally, the pattern introduced here should allow for a generic query design that supports different types of *control points*. These control points have a spatial–temporal position and associated attribute data [19]. In order to have a rigorous framework for the spatial–temporal calculations, yet remain flexible concerning the data representation, the space–time prism pattern presented here is based on Miller's [20] measurement theory for time geography and the semantic trajectories pattern introduced by Hu et al. [19]. Using a generic pattern lets applications hook into the pattern and then apply the queries defined at pattern level. The pattern therefore has to cover the elements that are required to calculate whether a space–time object is located inside a space–time prism or not.

4.2. Rationale and competency questions

For the purpose of data integration, it is sufficient to limit the space–time prism to the cases of single and double cones, as shown in Fig. 3. We limit this initial work to these cases for simplicity, though future work will attempt to extend this to more complex space–time concepts. We can thus reduce the pattern to the control point at its center, as well as the temporal (i.e., the tips of the cones) and spatial (i.e., the cone's diameter) extents. We limit the pattern to two-dimensional locations, both for simplicity and because three-dimensional location data are quite rare in the Web of Data.

In order to maximize flexibility in the application of the pattern, we do not restrict the thematic attributes of control points here. Likewise, we do not make any assumptions about *how* an attribute under consideration is related to the diameter of a cone. In classical time geography, that attribute would be maximum travel velocity; we discuss other examples in Section 5. We believe that both the kinds of thematic attributes and their relationship to the diameter of the cone are too application-specific to restrict here.

Based on this rationale, datasets modeled according to the proposed pattern will be able to answer the following questions. They have been phrased around application examples to better illustrate the use cases for the pattern:

Question 1. *Given a series of GPS trajectories, which control points are spatially–temporally close to a given control point, X?*

This question addresses the use case of obtaining control points from a predefined prism.

Question 2. *Given a series of measurements in a sensor network, which measurements are spatially–temporally related to the toxic plume that was released at control point X?*

This question covers the use case of inferring a prism based on the attribute values associated with the control points (i.e., when a dispersion rate is unknown).

Question 3. *Given an object O that has been observed at control point X_1 at time t_1, and at X_2 at t_2, what are the possible locations for O at t_3 (assuming constant speed)?*

Question 4. *Given n objects $O_1 \ldots O_n$ that have been observed at control points $X_1 \ldots X_n$ at times $t_1 \ldots t_n$, where and when can those objects meet (assuming constant speed)?*

The last two competency questions cover use cases from classical time geography, allowing us to answer questions about possible future locations of single objects and potential meeting points of multiple objects.

4.3. Ontology design pattern

Fig. 4 gives a schematic overview of the space–time prism pattern. It builds on elements defined in the semantic trajectories pattern [19] shown in gray; yet not all elements defined there are necessary to define the space–time prism. It is especially notable that the trajectory itself is not required for the data integration case we focus on here. While trajectories may have attributes that can serve as a basis for data integration, we focus on control points here and defer the integration of trajectory-based data through bundles (see Section 2) to future work.

For the formalization, we use the Web Ontology Language (OWL) and provide the axioms in Description Logic (DL) notation [47] both for readability and to be in line with the formalization of the trajectories pattern [19]. It makes use of the OWL-Time ontology [48] (prefix *time*) and the Points of Interest ontology design pattern[4] (prefix *poi*), which, when used in combination, allow for a flexible way to handle the temporal and spatial aspects of a control point. The design of our pattern is based on the definition of the space–time prism through:

1. the control point that the prism centers around (see Ax. (1a));
2. the points in time that define the tips of the cones (see Ax. (1b) and (1c));
3. the diameter of the prism, defining its spatial extent (see Ax. (1d)); and
4. the attribute that defines the diameter of the prism (see Ax. (2)).

The following DL axioms define the interplay between these elements:

$$Prism \sqsubseteq \leq 1 centeredAt.ControlPoint \tag{1a}$$
$$\sqcap \leq 1 time:hasBeginnnig.time:TemporalThing \tag{1b}$$
$$\sqcap \leq 1 time:hasEnd.time:TemporalThing \tag{1c}$$
$$\sqcap \leq 1 hasDiameter.diameter \tag{1d}$$

$$Diameter \sqsubseteq \exists defines^{-}.Attribute. \tag{2}$$

The attributes are deliberately undefined in order to maintain maximum flexibility concerning the attributes that can be used to

4 http://geog.ucsb.edu/~jano/POIpattern.png.

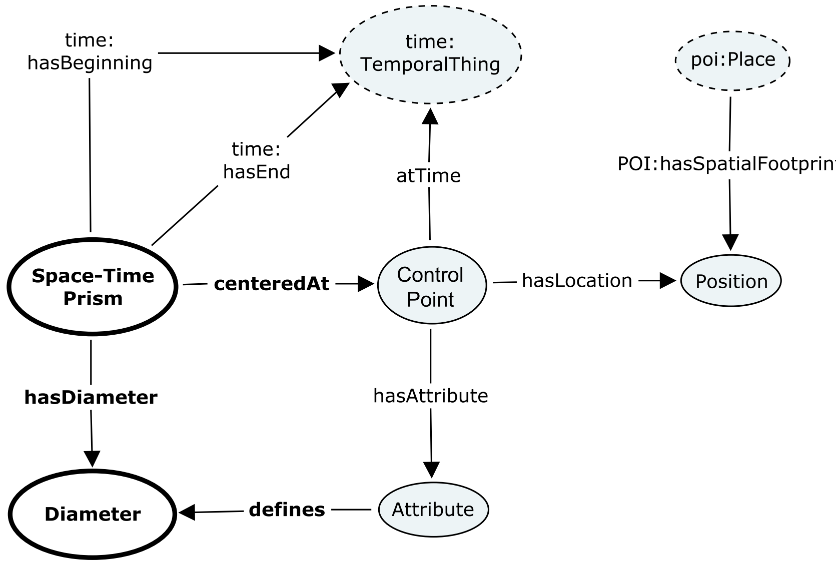

Fig. 4. Schematic overview of the space–time prism pattern.

define the diameter. For the control points, we adopt the notion introduced by Hu et al. [19] (they refer to the concept as a *fix*, which alludes to the language used in GPS positioning; we refer to the same concept as a control point here):

$$\begin{aligned} ControlPoint \sqsubseteq\ & \exists tTime.OWL - Time : Temporal\ Thing \\ & \sqcap \exists hasLocation.Position \\ & \sqcap \exists hasAttribute.Attribute \\ & \sqcap \exists hasFix^{-}.SemanticTrajectory. \end{aligned} \quad (3)$$

In classical time geography, the attribute that defines the diameter of the prism is maximum travel velocity; we generalize this to any observable phenomenon that is at least at ordinal scale to enable the calculation of spatial autocorrelation.[5] The separation of the beginning and end of the time span covered by the prism supports all cases where it is not desirable to have the control point in the middle of the time span. In the most extreme cases, this allows us to "search only backward" in a cone pointing back in time to look for control points that have influenced the control point under consideration. Likewise, we can "look only forward" (a cone pointing forward in time) to query for control points that have been influenced *by* the control point under consideration.

4.4. Integrated querying across different spatial–temporal models

The *Position* element (see Fig. 4) adopted from the semantic trajectories pattern is deliberately underspecified to support different encodings, such as the W3C Basic Geo Vocabulary for simple lat/lon points in WGS84,[6] or the GeoSPARQL ontology, building on the OGC simple features model to support any kind of vector geometry [51]. Likewise, it is possible to replace the control point in the pattern with different notions of spatial–temporal entities, such as events [52]. As long as it is possible to calculate Euclidean distance between different control points (i.e., distance in space–time), we can address the two cases of predefined and inferred prisms discussed in Section 3.2. Fig. 5 illustrates these two cases: If the beginning, end, and diameter of the prism (shown in blue) are known, we can calculate which other control points are spatial–temporally located within the prism (shown in red).

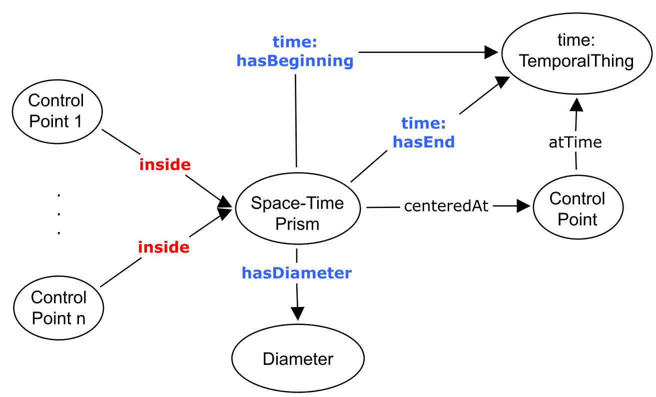

Fig. 5. Query pattern building on the space–time prism. The parts shown in red can be inferred from the parts shown in blue, and vice versa. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Conversely, if we know which control points are related, we can infer the spatial–temporal dimensions of the space–time prism.

By introducing the space–time prism as an additional element that is defined by the existing spatial and temporal dimensions specified by the respective pattern, the proposed pattern does not break other existing patterns, such as DOLCE's [37] participation pattern (see Fig. 2). The same goes for the setting pattern [53] that defines a spatial and temporal scoping for a setting. It has been developed with a focus on historical gazetteers, where the actual temporal scope of a period often depends on a spatial component (such as the *Bronze Age*, which is defined differently depending on region). In this case, the same scoping mechanism also works for the space–time prism, again introducing it as an additional element that can be used for querying and reconciliation.

5. Applications

The ontology design pattern introduced in the previous section can easily be turned into a *query pattern* by replacing parts of the pattern with variables. In the following, we introduce two applications for this query pattern, one from environmental modeling (Section 5.1) and one from cultural heritage (Section 5.2). In Section 5.3, we briefly discuss other potential applications.

5 While it may be possible to apply local indicators of spatial association for nominal scale data [49,50], we leave this for future work.

6 http://www.w3.org/2003/01/geo/.

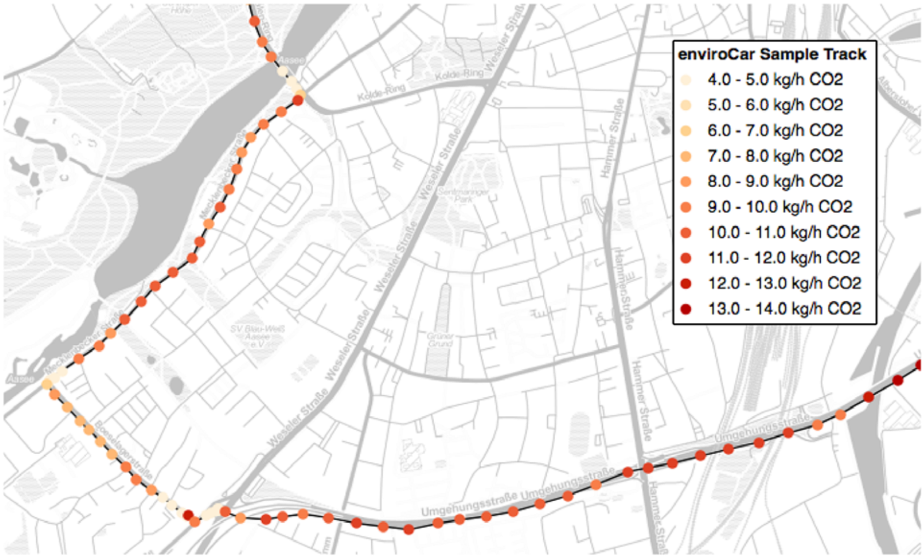

Fig. 6. Map showing a sample enviroCar trajectory, consisting of space–time control points and the connecting edges. Color scale of the control points represents the modeled CO_2 emissions. Additional variables, such as speed or number of revolutions per minute are also collected but not represented here. Basemap by Stamen Design, under CC BY 3.0. Data by OpenStreetMap, under ODbL. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5.1. EnviroCar

EnviroCar[7] is a community-based data collection platform for gathering vehicle-borne sensor data and producing environmental information [54]. EnviroCar uses standard Bluetooth OBD-II adaptors,[8] which are connected to a vehicle via the standard OBD connection that allows it to read parameters such as speed or revolutions per minute. From there, an Android smartphone records the data at regular time intervals, augmented with GPS information from the enviroCar smartphone app. The enviroCar app automatically calculates further information such as fuel consumption and CO_2 emissions, which can then be uploaded to the enviroCar platform server for subsequent analysis and sharing with the wider research and citizen science communities (see Fig. 6).

Understanding the spatial dynamics of fuel efficiency is a major area of research in environmental sustainability and road and vehicle efficiency studies. Similarly, the growing use of geo-sensor networks for air quality assessment (among other issues) requires integration of many disparate data sources, including mobile and static sensors. This provides an ideal use-case for the proposed space–time prism ontology design pattern. For example, suppose researchers working with the enviroCar data would like to determine whether a particularly high emissions reading is due to location (i.e., because there is a steep incline, intersection, or speed bump along the road) or largely location independent (i.e., simply due to vehicle type or driving style). Using the concept of a space–time prism presented in Section 3, it is possible to determine this.

For each control point in a series of enviroCar trajectories, we can compute the spatial autocorrelation between the control point and all other control points in the data. This determines the diameter of the space–time prism around each control point.[9] From here we get a relative measure of the degree of clustering in the data: very small space–time prisms indicate spatial effects (i.e., some environmental condition leading to increased emissions), whereas larger space–time prisms indicate a lack of spatial effects (i.e., possibly driving style). This result can then be used to validate said data by intersecting local environmental sensors from a sensor network with the space–time prisms generated from the enviroCar data. Observations can be *augmented* with sensor information: "show me all sensor measurements in the vicinity of an enviroCar emission peak", or *compared* with sensor observations: "show me all enviroCar control points that deviate strongly from local sensor network observations".

The following example is based on a measurement obtained from the enviroCar Linked Open Data API.[10] It shows a single control point – in this case, an observation based on the W3C Semantic Sensor Networks Ontology [56] – along with the measured CO_2 value at a given sampling time.

```
@base         <https://envirocar.org/api/stable/> .
@prefix geo:  <http://www.w3.org/2003/01/geo/wgs84_pos#> .
@prefix dul:  <http://www.loa-cnr.it/ontologies/DUL.owl#> .
@prefix ssn:  <http://purl.oclc.org/NET/ssnx/ssn#> .
@prefix xsd:  <http://www.w3.org/2001/XMLSchema#> .

  # this is a control point and its location in space:

<measurements/5207d871e4b058cd3d669c36> a ssn:Observation ;
  dul:hasPart <measurements/5207d871e4b058cd3d669c36#CO2> ;
  geo:lat "51.954"^^xsd:double ;
  geo:lon "7.65095"^^xsd:double .

  # the remaining triples define an attribute (CO2 output)
  # of the control point, and the control point's time stamp:

<measurements/5207d871e4b058cd3d669c36#CO2>
  ssn:observationResult
    <measurements/5207d871e4b058cd3d669c36#CO2_out> ;
  ssn:observationSamplingTime
    <measurements/5207d871e4b058cd3d669c36#time> ;
  ssn:observedProperty <phenomenons/CO2> .
```

7 https://www.envirocar.org.

8 http://www.obdii.com/background.html.

9 This computation can be accomplished in several ways, such as by using AMOEBA [55].

10 See http://envirocar.github.io/enviroCar-server/lod/.

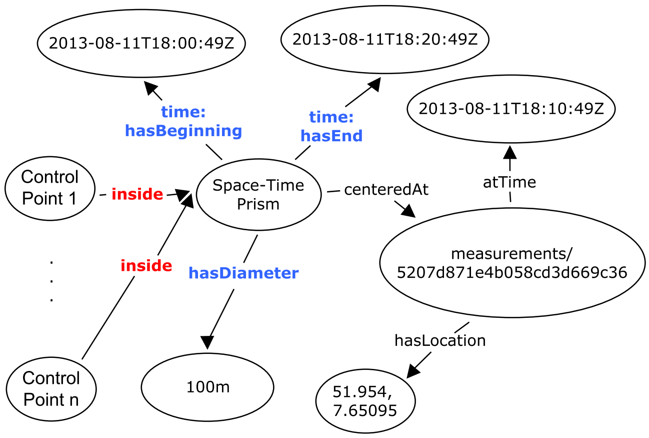

Fig. 7. Predefined space–time prism query pattern for the enviroCar example. With the properties set (shown in blue), we can calculate which other control points are inside the prism (shown in red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

```
<measurements/5207d871e4b058cd3d669c36#CO2_out>
  a ssn:SensorOutput ;
  ssn:hasValue
    <measurements/5207d871e4b058cd3d669c36#CO2_value> .

<measurements/5207d871e4b058cd3d669c36#CO2_value> a dul:Amount ;
  dul:hasDataValue "0.004162327586206897"^^xsd:double ;
  dul:isClassifiedBy <phenomenons/CO2#unit> .

<measurements/5207d871e4b058cd3d669c36#time> a dul:TimeInterval ;
  dul:hasDataValue "2013-08-11T18:10:49Z"^^xsd:dateTime .
```

Taking the dispersion rate of CO_2 into account, we can match the query pattern for the predefined space–time prism from Section 4.4 to this dataset, as shown in Fig. 7. In this example, we are temporally searching the 10 min before and after the control point respectively, and set the diameter for the prism to 100 m:

```
base <https://envirocar.org/api/stable/>
prefix geo: <http://www.opengis.net/ont/geosparql#>
prefix geof: <http://www.opengis.net/def/function/geosparql/>
prefix dul: <http://www.loa-cnr.it/ontologies/DUL.owl#>
prefix ssn: <http://purl.oclc.org/NET/ssnx/ssn#>
prefix xsd: <http://www.w3.org/2001/XMLSchema#>
prefix units: <http://www.opengis.net/def/uom/OGC/1.0/>

SELECT * WHERE {

  # specify the parameters for this search:

  BIND ("600"^^xsd:integer AS ?duration) . # duration (in secs)
  BIND ("100"^^xsd:integer AS ?diameter) . # prism diameter (in m)

  # this is the center of our prism

  <measurements/cp1> geo:hasGeometry ?geometry1 ;
  dul:hasPart ?part1 .

  # fetch its location

  ?geometry1 geo:asWkt ?wkt1 .

  # fetch the timestamp for the CO2 observation at the control point

  ?part1 ssn:observationSamplingTime ?timeInstance1 .
  ?timeInstance1 dul:hasDataValue ?ts1 .

  #get the same data for other control points

  ?cp geo:hasGeometry ?geometry2 ;
  dul:hasPart ?part2 .

  ?geometry2 geo:asWkt ?wkt2 .

  ?part2 ssn:observationSamplingTime ?timeInstance2 .
  ?timeInstance2 dul:hasDataValue ?ts2 .

  # we make the simplifying assumption here that both samplings
  have occurred on the same day and in the same time zone. This can
  relatively easily be extended to any arbitrary sampling times

  BIND (((hours(?ts1) * 3600) + (minutes(?ts1) * 60) +
           xsd:integer(seconds(?ts1))) AS ?secsOfDay1).
  BIND (((hours(?ts2) * 3600) + (minutes(?ts2) * 60) +
           xsd:integer(seconds(?ts2))) AS ?secsOfDay2).

  # specify the "tips" of the cones — 10 minutes (= 600 secs) into the
  past from ?secsOfDay1

  BIND ((?secsOfDay1 - ?duration) AS ?prismStart) .

  #...and 10 minutes in the future:

  BIND ((?secsOfDay1 + ?duration) AS ?prismEnd) .

  # figure out at which point of the timespan we are to calculate the
  diameter of the prism at this point in time ("disc")

  BIND ((abs(?secsOfDay1 - ?secsOfDay2)) / ?duration AS ?section) .

  # calculate disc diameter

  BIND (?diameter * ?section AS ?discDiameter) .

  # filter the results to those within the timespan of the prism

  FILTER (?secsOfDay2 > ?prismStart) .
  FILTER (?secsOfDay2 < ?prismEnd) .

  # filter results to the locations that are within the "disc"

  FILTER(geof:distance(?wkt1, ?wkt2, units:metre) < ?discDiameter)
}
```

The query strategy consists in (1) querying all control points that fall within the given time frame; (2) limit this selection to all control points that fall within the cylinder defined by the diameter of the space–time prism; and (3) for each remaining candidate, calculate the disc slice [20] of the prism at the given point in time and check whether the candidate control point is inside the disc. The enviroCar example shows that the intuitive notion of a space–time prism as defined in Section 4 can easily be turned into a query that yields the correct results, i.e., it returns all control points that are spatially and temporally contained in the 3D cone defined by the dispersion rate of CO_2.

5.2. Federated queries for cultural heritage data

The enviroCar use case shows that space–time prisms can be used to query the Web of Data with the technologies underlying the Linked Data approach. While this is certainly convenient, the combination of spatial–temporal reasoning with semantic reasoning (e.g., subsumption reasoning) *can* already be

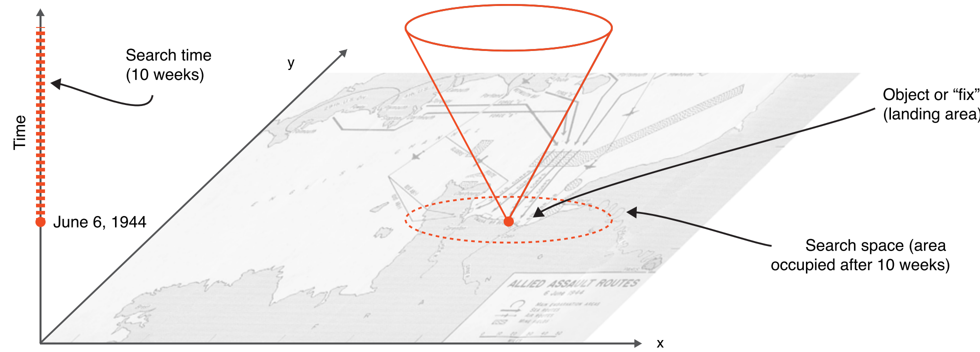

Fig. 8. Forward-facing space–time cone approximating the progress of the Allied forces in the 10 weeks after D-Day.

accomplished using existing technologies (e.g., GeoSPARQL, W3C Geo, and OWL-Time, etc.). The strength of this approach, therefore, comes from the *direct* application of intuitive time geography concepts to *distributed* datasets.

In order to further illustrate the capabilities of our proposed design pattern, we introduce a use case that applies a query based on space–time prisms to cultural heritage data. More specifically, we are looking for photos that were taken during the first few weeks after *D-Day* – the invasion of Allied forces in Normandy, France – on June 6, 1944. In the weeks following D-Day, the Allied forces gained control over a coastal zone of about 100 miles in diameter, which ultimately led to the restoration of the French Republic and was a major contribution to the Allied victory in World War II. In terms of time geography, our query is based on a forward-facing cone, with the tip set to the landing site in Normandy, a duration of 10 weeks, and a diameter of 100 miles (see Fig. 8).

Retrieving photos inside this cone requires information from the following Linked Data sources:

- DBpedia[11] provides information about D-Day, including the date the operation started and coordinates for the location of the landing.
- Europeana[12] [1] provides metadata about cultural heritage artifacts from a large number of European museums, including the photos we are looking for.
- GeoNames[13] provides coordinates and the administrative hierarchy for all places referenced from the Europeana data.
- GEMET,[14] the GEneral Multilingual Environmental Thesaurus, provides the classification scheme for artifacts in Europeana.

GEMET not only provides us with the correct identifier for photographs,[15] but also with the corresponding subsumption hierarchy. The `skos:broader` and `skos:narrower` relationships allow us to include aerial images in our query, which are a narrower type of photograph in GEMET. Likewise, we can include *all* narrower types of photographs at the same time using SPARQL 1.1 property paths [57]:

```
prefix skos: <http://www.w3.org/2004/02/skos/core#>
prefix gemet: <http://www.eionet.europa.eu/gemet/concept/>
prefix edm: <http://www.europeana.eu/schemas/edm/>

SELECT ?artifact WHERE {
    ?artifact edm:hasMet ?category .
    ?category skos:narrower* gemet:6205 .
}
```

The following query demonstrates how to combine results from different SPARQL endpoints in a federated query. It selects all places from a local store that are within 10 miles of the landing location of the invasion of Normandy, retrieved from DBpedia.

```
prefix wgs84: <http://www.w3.org/2003/01/geo/wgs84_pos#>
prefix dpbo: <http://dbpedia.org/ontology/>
prefix geof: <http://www.opengis.net/def/function/geosparql/>

SELECT * WHERE {

    ?place wgs84:geometry ?location .

    SERVICE <http://dbpedia.org/sparql> {
        <http://dbpedia.org/resource/Invasion_of_Normandy>
            wgs84:geometry ?landing ;
            dpbo:date ?date . }

   FILTER(geof:distance(?landing, ?location, units:mile) < "10")
}
```

Combining such federated queries and semantic reasoning with the space–time prism query pattern introduced in Section 5.1 would follow the same approach as in the previous section: first filter temporally, then filter spatially, and for each of the remaining candidates, calculate the slice radius at its position in the cone to determine whether it is inside.[16] This example demonstrates that the intuitive notion of space–time prisms and the corresponding pattern enable powerful spatial–temporal semantic information retrieval across distributed data sources.

The main limiting factor for this kind of query is data availability. The particular example illustrated above, for example, does not bear any meaningful results on the listed data sources yet, since the spatial and temporal granularity of information in the datasets is not high enough. Photographs in Europeana are only annotated with the year they were taken, and mostly do not contain information about the location at which they were taken. Moreover, some of the datasets that contain relevant data for such queries, such as the Library of Congress Linked Data Service[17] do not yet offer a SPARQL endpoint.

5.3. Other applications

The two examples discussed above illustrate the utility of our approach for environmental monitoring and retrieval of artifacts. However, a number of applications in other – and potentially very different – domains can benefit from this approach. One such

11 See http://dbpedia.org.
12 See http://europeana.eu. SPARQL endpoint available at http://europeana.ontotext.com/sparql.
13 See http://geonames.org.
14 See http://www.eionet.europa.eu/gemet/.
15 See http://www.eionet.europa.eu/gemet/concept/6205.
16 See http://carsten.io/photoquery.sparql for a complete example.
17 http://id.loc.gov.

example is the spatial@linkedscience[18] effort, which provides an access point to all papers ever published in the major Geographic Information conferences [58]. In this context, the pattern could be used to track the reach of a particular research group through space and time, with the publications as control points and measures of topic similarity as the spatial autocorrelation metric. By examining the changing shape of the space–time prisms, we can get an idea of the spread rate and reach of particular localities (universities).

In previous work, we have shown how to semantically reason over historic map contents using Linked Data [59]. Adding time geography based queries to this approach would allow us to query for artifacts that potentially contain information about persons or events based on documented appearances. Other examples include the collection of publicly available pictures taken around a given event, the historical events that have influenced a book or a piece of art, and potentially even the identification of potential witnesses after a crime based on cell phone locations.

6. Conclusions

In this paper, we presented an ontology design pattern for space–time prisms, a concept from time geography, and demonstrated its utility using two real-world examples. The pattern is based on Miller's measurement theory for time geography [20] and facilitates querying across, and integration of, disparate data sources through a re-conceptualization of the space–time prism. Building on previous work on semantic trajectories [19], the pattern is formalized in the Web Ontology Language to facilitate its application with data following Linked Data principles. While other technological approaches may yield the same or similar results, the pattern enables the use of an intuitive conceptual framework in an existing and well-established technology stack.

Time geography provides a rich array of concepts, models, and rigorous analytical definitions of spatial–temporal entities with which to understand space–time interactions. We have demonstrated that the space–time prism is a viable tool to explore the relationships within and between Linked Data sources. Our approach goes beyond previous work on time geography as it shows that the space–time prism can be used as a tool for data integration and exploration, while at the same time leveraging the decentralized semantic reasoning capabilities of the Linked Data approach.

One possible extension to this framework is the use of spatial cross-correlation measures to define *inferred* space–time prism diameters. This would allow queries *across* different data types to be integrated within the approach presented here. Examples where this might be useful include questions such as "where are high CO_2 emissions associated with infant mortality rates?", or "at what temporal lag do predator and prey cycles operate?".

References

[1] A. Isaac, B. Haslhofer, Europeana linked open data—data.europeana.eu, Semant. Web 4 (3) (2013) 291–297.

[2] M. Samwald, A. Jentzsch, C. Bouton, C. Kallesøe, E. Willighagen, J. Hajagos, M. Marshall, E. Prud'hommeaux, O. Hassanzadeh, E. Pichler, et al., Linked open drug data for pharmaceutical research and development, J. Cheminformatics 3 (1) (2011) 19.

[3] C. Keßler, C. Hendrix, The humanitarian exchange language: Coordinating disaster response with semantic web technologies, Semant. Web J. 6 (1) (2015) 6–21. http://dx.doi.org/10.3233/SW-1301304.

[4] H. Patni, C. Henson, A. Sheth, Linked sensor data, in: Collaborative Technologies and Systems, IEEE CTS 2010, 2010, pp. 362–370.

[5] T. Berners-Lee, Linked Data—Design Issues, 2009. Online: http://www.w3.org/DesignIssues/LinkedData.html.

[6] G. Kobilarov, T. Scott, Y. Raimond, S. Oliver, C. Sizemore, M. Smethurst, C. Bizer, R. Lee, Media meets semantic web—How the BBC uses DBpedia and linked data to make connections, in: The Semantic Web: Research and Applications, Springer, 2009, pp. 723–737.

[7] A.J.G. Gray, P. Groth, A. Loizou, S. Askjaer, C. Brenninkmeijer, K. Burger, C. Chichester, C.T. Evelo, C. Goble, L. Harland, S. Pettifer, M. Thompson, A. Waagmeester, A.J. Williams, Applying linked data approaches to pharmacology: Architectural decisions and implementation, Semant. Web 5 (2) (2014) 101–113.

[8] M. d'Aquin, Putting linked data to use in a large higher-education organisation, in: Workshop: Interacting with Linked Data at Extended Semantic Web Conference, ESWC 2012, 2012.

[9] C. Keßler, T. Kauppinen, Linked open data university of Münster—Infrastructure and applications, in: E. Simperl, B. Norton, D. Mladenic, E. Della Valle, I. Fundulaki, A. Passant, R. Troncy (Eds.), The Semantic Web: ESWC 2012 Satellite Events, in: Lecture Notes in Computer Science, vol. 7540, Springer, Berlin, Heidelberg, 2015, pp. 447–451.

[10] T. Kauppinen, G.M. de Espindola, J. Jones, A. Sánchez, B. Gräler, T. Bartoschek, Linked Brazilian Amazon rainforest data, Semant. Web 5 (2) (2014) 151–155.

[11] T. Hägerstrand, What about people in regional science? Pap. Reg. Sci. Assoc. 24 (1970) 7–21.

[12] M. Raubal, H.J. Miller, S. Bridwell, User-centred time geography for location-based services, Geogr. Ann.: Ser. B Hum. Geogr. 86 (4) (2004) 245–265.

[13] P. Crease, T. Reichenbacher, Linking time geography and activity theory to support the activities of mobile information seekers, Trans. GIS 17 (4) (2013) 507–525.

[14] M. Espeter, M. Raubal, Location-based decision support for user groups, J. Locat. Based Serv. 3 (3) (2009) 165–187.

[15] S. Winter, M. Raubal, Time geography for ad-hoc shared-ride trip planning, in: 7th International Conference on Mobile Data Management, 2006. MDM 2006, 2006.

[16] M. Raubal, S. Winter, S. Teßmann, C. Gaisbauer, Time geography for ad-hoc shared-ride trip planning in mobile geosensor networks, ISPRS J. Photogramm. Remote Sens. 62 (5) (2007) 366–381.

[17] P. Gärdenfors, Conceptual Spaces—The Geometry of Thought, Bradford Books, MIT Press, Cambridge, MA, 2000.

[18] A. Gangemi, V. Presutti, Ontology design patterns, in: Handbook on Ontologies, Springer, 2009, pp. 221–243.

[19] Y. Hu, K. Janowicz, D. Carral, S. Scheider, W. Kuhn, G. Berg-Cross, P. Hitzler, M. Dean, D. Kolas, A geo-ontology design pattern for semantic trajectories, in: Spatial Information Theory, Springer, 2013, pp. 438–456.

[20] H.J. Miller, A measurement theory for time geography, Geogr. Anal. 37 (1) (2005) 17–45.

[21] H.J. Miller, What about people in geographic information science?, 2005.

[22] S.-L. Shaw, Guest editorial introduction: time geography—its past, present and future, J. Transp. Geogr. 23 (0) (2012) 1–4. Special Issue on Time Geography.

[23] A. Moore, P. Whigham, A. Holt, C. Aldridge, K. Hodge, A time geography approach to the visualisation of sport, in: Proceedings of the 7th International Conference on GeoComputation, University of Southampton, United Kingdom, September 8th–10th, 2003.

[24] J.A. Long, T.A. Nelson, F.S. Nathoo, Toward a kinetic-based probabilistic time geography, Int. J. Geogr. Inf. Sci. 28 (5) (2014) 855–874.

[25] H.J. Miller, Modelling accessibility using space–time prism concepts within geographical information systems, Int. J. Geogr. Inf. Syst. 5 (3) (1991) 287–301.

[26] J.A. Long, T.A. Nelson, Time geography and wildlife home range delineation, J. Wildl. Manage. 76 (2) (2012) 407–413.

[27] M.-P. Kwan, Gender and individual access to urban opportunities: A study using space–time measures, Prof. Geogr. 51 (2) (1999) 210–227.

[28] D. Rainham, I. McDowell, D. Krewski, M. Sawada, Conceptualizing the healthscape: Contributions of time geography, location technologies and spatial ecology to place and health research, Soc. Sci. Med. 70 (5) (2010) 668–676.

[29] M.-P. Kwan, GIS methods in time-geographic research: Geocomputation and geovisualization of human activity patterns, Geogr. Ann.: Ser. B Hum. Geogr. 86 (4) (2004) 267–280.

[30] M. Raubal, Representing concepts in time, in: C. Freksa, N. Newcombe, P. Gärdenfors, S. Wölfl (Eds.), Spatial Cognition VI. Learning, Reasoning, and Talking About Space, Springer, 2008, pp. 328–343.

[31] S. Winter, Towards a probabilistic time geography, in: Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, November 4–6, 2009, Seatle, WA, ACM, 2009, pp. 528–531.

[32] S. Winter, Z.-C. Yin, The elements of probabilistic time geography, GeoInformatica 15 (3) (2011) 417–434.

[33] H.J. Miller, S.A. Bridwell, A field-based theory for time geography, Ann. Assoc. Am. Geogr. 99 (1) (2009) 49–75.

[34] E. Gamma, R. Helm, R. Johnson, J. Vlissides, Design Patterns: Elements of Reusable Object-Oriented Software, Pearson Education, 1994.

[35] A. Gangemi, V. Presutti, Towards a pattern science for the semantic, Semant. Web 1 (1) (2010) 61–68.

[36] A. Gangemi, Ontology design patterns for semantic web content, in: Y. Gil, E. Motta, V.R. Benjamins, M.A. Musen (Eds.), The Semantic Web–ISWC 2005, Springer, 2005, pp. 262–276.

[37] C. Masolo, A. Gangemi, N. Guarino, A. Oltramari, L. Schneider, WonderWeb deliverable D18: The wonderweb library of foundational ontologies, 2004.

[38] A. Degbelo, An ontology design pattern for spatial data quality characterization in the semantic sensor web, in: O. C. Cory Henson, Kerry Taylor, (Eds.), Proceedings of the 5th International Workshop on Semantic Sensor Networks, 11th International Semantic Web Conference 2012, ISWC 2012, Boston, Massachusetts, USA, November 12, 2012, 2012.

[18] http://spatial.linkedscience.org/.

[39] D.C. Martínez, K. Janowicz, P. Hitzler, A logical geo-ontology design pattern for quantifying over types, in: Proceedings of the 20th International Conference on Advances in Geographic Information Systems, SIGSPATIAL'12, ACM, New York, NY, USA, 2012, pp. 239–248.
[40] D. Carral, S. Scheider, K. Janowicz, C. Vardeman, A. Krisnadhi, P. Hitzler, An ontology design pattern for cartographic map scaling, in: P. Cimiano, O. Corcho, V. Presutti, L. Hollink, S. Rudolph (Eds.), The Semantic Web: Semantics and Big Data, in: Lecture Notes in Computer Science, vol. 7882, Springer, Berlin, Heidelberg, 2013, pp. 76–93.
[41] K. Janowicz, M. Compton, K. Taylor, A. Ayyagari, D.D. Roure, The stimulus-sensor-observation ontology design pattern and its integration into the semantic sensor network ontology, in: The 3rd International Workshop on Semantic Sensor Networks 2010 (SSN10) in Conjunction with the 9th International Semantic Web Conference, ISWC 2010, 2010.
[42] M. Baglioni, J. Macedo, C. Renso, M. Wachowicz, An ontology-based approach for the semantic modelling and reasoning on trajectories, in: Il-Yeol Song, et al. (Eds.), Advances in Conceptual Modeling—Challenges and Opportunities, in: Lecture Notes in Computer Science, vol. 5232, Springer, 2008, pp. 344–353.
[43] D. Zheni, A. Frihida, H.B. Ghezala, C. Claramunt, A semantic approach for the modeling of trajectories in space and time, in: Advances in Conceptual Modeling-Challenging Perspectives, Springer, 2009, pp. 347–356.
[44] W.R. Tobler, A computer movie simulating urban growth in the detroit region, Econ. Geogr. (1970) 234–240.
[45] B. Boots, Local measures of spatial association, Ecoscience 9 (2002) 168–176.
[46] L. Anselin, Local indicators of spatial association–LISA, Geogr. Anal. 27 (2) (1995) 93–115.
[47] Hitzler P, Krotzsch M, Rudolph S, Foundations of Semantic Web Technologies, CRC Press, 2011.
[48] J.R. Hobbs, F. Pan, Time ontology in OWL, 2006. W3C Working Draft: http://www.w3.org/TR/owl-time.
[49] B. Boots, Developing local measures of spatial association for categorical data, J. Geogr. Syst. 5 (2) (2003) 139–160.
[50] J.A. Long, T.A. Nelson, M.A. Wulder, Local indicators for categorical data: impacts of scaling decisions, Can. Geogr. / Géogr. Can. 54 (1) (2010) 15–28.
[51] Open Geospatial Consortium. OGC GeoSPARQL—A Geographic Query Language for RDF Data, 2012. Online: http://opengis.org/standards/geosparql.
[52] J. Trame, C. Keßler, W. Kuhn, Linked data and time—Modeling researcher life lines by events, in: T. Tenbrink, J. Stell, A. Galton, Z. Wood (Eds.), Spatial Information Theory, in: Lecture Notes in Computer Science, vol. 8116, Springer, 2013, pp. 205–223.
[53] K. Grossner, K. Janowicz, C. Keßler, Place, period, and setting for linked data gazetteers, in: R. Mostern, H. Southall, M. Berman (Eds.), Placing Names: Enriching and Integrating Gazetteers, Indiana University Press, Bloomington, 2015, in press.
[54] A. Bröring, A. Remke, C. Stasch, C. Autermann, M. Rieke, J. Möllers, enviroCar: A citizen science platform for analyzing and mapping crowd-sourced car sensor data, Trans. GIS 19 (3) (2015) 362–376.
[55] J. Aldstadt, A. Getis, Using AMOEBA to create a spatial weights matrix and identify spatial clusters, Geogr. Anal. 38 (4) (2006) 327–343.
[56] M. Compton, P. Barnaghi, L. Bermudez, R. García-Castro, O. Corcho, S. Cox, J. Graybeal, M. Hauswirth, C. Henson, A. Herzog, et al., The SSN ontology of the W3C semantic sensor network incubator group, J. Web Semant. 17 (2012) 25–32.
[57] S. Harris, A. Seaborne, SPARQL 1.1 Query Language, 2013. W3C recommendation available from: http://www.w3.org/TR/sparql11-query/.
[58] C. Keßler, K. Janowicz, T. Kauppinen, spatial@linkedscience–Exploring the research field of GIScience with linked data, in: N. Xiao, M.-P. Kwan, M.F. Goodchild, S. Shekhar (Eds.), GIScience 2012, Columbus, OH, USA, September 18–21, 2012, in: Lecture Notes in Computer Science, vol. 7478, Springer, 2012, pp. 102–115.
[59] S. Scheider, J. Jones, A. Sánchez, C. Keßler, Encoding and querying historic map content, in: J. Huerta, S. Schade, C. Granell (Eds.), Connecting a Digital Europe Through Location and Place, in: Lecture Notes in Geoinformation and Cartography., Springer International Publishing, 2014, pp. 251–273.