Best Practices for Publishing, Retrieving, and Using Spatial Data on the Web

Linda van den Brink, Payam Barnaghi, Jeremy Tandy, Ghislain Atemezing, Rob Atkinson, Byron Cochrane, Yasmin Fathy, Raúl García Castro, Armin Haller, Andreas Harth, et al.

► To cite this version:

Linda van den Brink, Payam Barnaghi, Jeremy Tandy, Ghislain Atemezing, Rob Atkinson, et al.. Best Practices for Publishing, Retrieving, and Using Spatial Data on the Web. Semantic Web – Interoperability, Usability, Applicability, In press. hal-01885342

HAL Id: hal-01885342

<https://hal.science/hal-01885342>

Submitted on 1 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Best Practices for Publishing, Retrieving, and Using Spatial Data on the Web

Linda van den Brink ^{a,*}, Payam Barnaghi ^b, Jeremy Tandy ^c, Ghislain Atemezang ^d, Rob Atkinson ^e, Byron Cochrane ^f, Yasmin Fathy ^b, Raúl García Castro ^g, Armin Haller ^h, Andreas Harth ⁱ, Krzysztof Janowicz ^j, Şefki Kolozali ^k, Bart van Leeuwen ^l, Maxime Lefrançois ^m, Josh Lieberman ⁿ, Andrea Perego ^{o,**}, Danh Le-Phuoc ^p, Bill Roberts ^q, Kerry Taylor ^h and Raphaël Troncy ^r

^a *Geonovum, Barchman Wuytierslaan 10, Postbus 508, 3800 AM, Amersfoort, The Netherlands*

^b *Institute for Communication Systems, University of Surrey, United Kingdom*

^c *MetOffice, United Kingdom*

^d *Mondeca, France*

^e *Metalinkage, Australia*

^f *Land Information New Zealand, New Zealand*

^g *Computer Science School, Universidad Politécnica de Madrid, Spain*

^h *College of Engineering & Computer Science, Australian National University, Australia*

ⁱ *Institute AIFB, Karlsruhe Institute of Technology, Germany*

^j *Geography Department, University of California, Santa Barbara, CA, USA*

^k *Analytical & Environmental Sciences Department, King's College London, United Kingdom*

^l *Netage, The Netherlands*

^m *École des Mines de Saint-Étienne, France*

ⁿ *Center for Geographic Analysis, Harvard University, MA, USA*

^o *Joint Research Centre, European Commission, Italy*

^p *Open Distributed Systems, Technische Universität Berlin, Germany*

^q *Swirrl, United Kingdom*

^r *Data Science Department, Eurecom, France*

Editor: Pascal Hitzler, Wright State University, USA

Solicited reviews: Xiaogang Ma, University of Idaho, USA; Christoph Schlieder, Universität Bamberg, Germany; two anonymous reviewers

Abstract. Data owners are creating an ever richer set of information resources online, and these are being used for more and more applications. Spatial data on the Web is becoming ubiquitous and voluminous with the rapid growth of location-based services, spatial technologies, dynamic location-based data and services published by different organizations. However, the heterogeneity and the peculiarities of spatial data, such as the use of different coordinate reference systems, make it difficult for data users, Web applications, and services to discover, interpret and use the information in the large and distributed system that is the Web. To make spatial data more effectively available, this paper summarizes the work of the joint W3C/OGC Working Group on Spatial Data on the Web that identifies 14 best practices for publishing spatial data on the Web. The paper extends that work by presenting the identified challenges and rationale for selection of the recommended best practices, framed by the set of principles that guided the selection. It describes best practices that are employed to enable publishing, discovery and retrieving (querying) spatial data on the Web, and identifies some areas where a best practice has not yet emerged.

Keywords: Geographic information systems, Spatial data, Web technologies, World Wide Web, W3C, Open Geospatial Consortium, OGC

*Corresponding author. E-mail: l.vandenbrink@geonovum.nl.

**The views expressed are purely those of the author and may not in any circumstances be regarded as stating an official position of the European Commission.

1570-0844/0-1900/\$35.00 © 0 – IOS Press and the authors. All rights reserved

1. Introduction

Spatial data is important. Firstly, because it has become ubiquitous with the explosive growth in positioning technologies attached to mobile vehicles, portable devices, and autonomous systems. Secondly, because it is fundamentally useful for countless convenient consumer services like transport planning, or for solving the biggest global challenges like climate change adaptation [1]. Historically, sourcing, managing and using high-quality spatial data has largely been the preserve of military, government and scientific enterprises. These groups have long recognized the importance and value that can be obtained by sharing their own specialized data with others to achieve cross-theme interoperability, increased usability and better spatial awareness, but they have struggled to achieve the cross-community uptake they would like. Spatial Data Infrastructures (SDIs) [2], which commonly employ the mature representation and access standards of the Open Geospatial Consortium (OGC), are now well developed, but have become a part of the “deep Web” that is hidden for most Web search engines and human information-seekers. Even geospatial experts still do not know where to start looking for what they need or how to use it when they find it. The integration of spatial data from different sources offers possibilities to infer and gain new information; however, spatial data on the Web is published in various structures, formats and with different granularities. This makes publishing, discovering, retrieving, and interpreting the spatial data on the Web a challenging task. By contrast, the linked data Web, as a platform of principles, tools, and standards championed by the World Wide Web Consortium (W3C) enables data discoverability and usability that is readily visible in, for example, search engine results for consumer shopping. The principles are based on proven aspects of the Web such as resolvable identifiers, common representation formats, and rich interlinking of independently-published information, but adds explicit vocabulary management and tooling that targets the huge Web developer community. Can these principles be successfully applied to the world of complex spatial data to achieve the desired usability and utility?

There are, already, many good examples of projects and Web services that deliver to these goals, such as spatial data publication platforms in the Netherlands,¹

Nanaimo City in Canada,² or the UK Environmental Agency,³ but also popular spatial data collections such as Geonames.⁴ However, spatial data custodians struggle to find the *best* way to publish their data in order to optimize the future impact as more data appears, more tools are developed, and the consumer community grows. Similarly, Web developers as data consumers and tool developers as foundation-stonemasons are demanding an expert consensus to guide their product development. The W3C and OGC standardization bodies jointly convened a large workshop in London in 2014 where these issues were extensively discussed over two days⁵. As a result of the interest, enthusiasm and challenges identified there, they proceeded to establish a joint working group to develop, amongst other things, a compendium of best practices for spatial data on the Web, published in September 2017 [3]. The working group completed its work in 2017, but was succeeded by an interest group⁶ in which the work is continued.

This paper is a companion publication to the Spatial Data on the Web Best Practices Note that was endorsed by the working group, including the authors of this paper. It summarizes the work and describes the best practices themselves, but additionally presents the principles that guided the selection of these best practices, and the rationale underlying particular selections. It also identifies some areas where a best practice seems to be needed but has not yet emerged.

1.1. Background: spatial data, the Web, and semantics

Any data (or metadata) that has a location component can be viewed as spatial data: its spatial nature means certain operations such as proximity and containment functions have a meaning within the spatial domain. A location component is a reference to a place on Earth or within some other space (e.g., another planet, or a shopping mall) and can be many things: a physical object with a fixed location, such as a building or canal; an administrative unit, like a municipality or postal code area, or the trajectory of a moving object like a car. The power of spatial data is in the oppor-

²<http://maps.nanaimo.ca/data/>

³<http://environment.data.gov.uk/bwq/profiles/>

⁴<http://www.geonames.org>

⁵Linking Geospatial Data (LGD'14). 5-6 March 2014, London. <https://www.w3.org/2014/03/lgd/>

⁶<https://www.w3.org/2017/sdwig/charter.html>

¹<https://data.pdok.nl/datasets>

Table 1
Standardized aspects of SDIs

Aspect	Description	Reference standards
Discoverability	Annotate resources with metadata	ISO 19115 [4, 5], ISO 19119 [6, 7]
Accessibility	Web services for discovering, viewing, downloading, sharing geospatial raster data (coverages) etc.	OGC CSW [8], OGC WMS [9], OGC WFS [10], OGC WCS [11]
Portrayal	Defining rules for displaying spatial data	ISO 19117 [12], OGC SLD [13], OGC SE [14], OGC KML [15]
Information modeling	Describing the contents of information resources, including geometry	ISO 19103 [16], ISO 19107 [17], ISO 19109 [18], ISO 19110 [19]
Data exchange	Defining formats for exchanging the data	OGC GML [20]
Spatial reference systems	Specifying the location on Earth of geographical information	ISO 19111 [21, 22]

tunity to combine and integrate information based on location. While spatial data refers to any data that has a location component either on Earth or within some other space such as another planet (as defined earlier), geospatial data refers to any data that has a location in terms of geographic coordinates (e.g., GIS coordinates) within Earth. The focus of the Spatial Data on the Web Best Practices, and consequently of this paper, is on geospatial data. Although non-geospatial use cases were brought to the working group and were included in the Spatial Data on the Web Use Cases and Requirements (UCR) [23], there were no active participants in the working group who had expertise with spatial data other than geospatial. The focus was therefore narrowed to geospatial data; requirements of non-geospatial data might be included in future work. This paper likewise deals almost exclusively with geospatial data. That said, many of the best practices are applicable to wider spatial data concerns. In the remainder of this paper, we simply refer to spatial data for brevity.

Spatial data can also have a temporal dimension. Temporal data varies over time. On the other hand, spatio-temporal data captures spatial data (i.e. current location) associated with the time the data is taken at that particular location [24]. However, in order to use spatial-temporal data, it has to be available and accessible. Common practice is to publish this data using an SDI. SDIs are based on a service-oriented architecture (SOA), in which existing resources are documented using dataset-level metadata, published in catalogs which are the most important discovery and access mechanism [25]. More detailed metadata describing structure and content of datasets, as well as service requests and payloads, are far less commonly shared, and whilst standards suitable for some aspects of this

requirement are emerging, defining a common or best practice remains a challenge.

The Open Geospatial Consortium (OGC), founded in 1994, publishes technical standards necessary for SDIs to work in an interoperable way. These standards are based on the more abstract standards for geospatial information from the ISO Technical Committee 211 on Geographic information / Geomatics (ISO/TC 211). Different aspects of the SDI are standardized (see Table 1).

The OGC developed the first standards for spatial data Web services as of 1998 and has responded to early architectural trends in the Web (e.g., SOA). While SDIs and related standards were developing, so did the World Wide Web. Web standards like HTML, XML and RDF [26] were created in the nineties as well. While the Web started off as mostly documents with hyperlinks, over the years it evolved to much more sophisticated Web applications, including mass applications in which geospatial data was used, like Bing Maps, Google Maps, Google Earth and OpenStreetMap. More recently, XML has been replaced in many Web applications by more lightweight formats (e.g. JSON, and RDF).

As a generic model or framework, RDF can be used to publish geographic information. Its strengths include its structural flexibility, particularly suited for rich and varied forms for metadata required for different purposes. However, it has no specific features for encoding geometry, which is central to geographic information. Several vocabularies and extensions have been proposed for this purpose, including a core RDF/OWL vocabulary for geographic information which is part of OGC's GeoSPARQL [27].

Figure 1 provides an illustration of the main areas of overlap between the different communities of practice of the Spatial, Semantic and the “mainstream” Web—

Fig. 1. Commonalities between different communities of practice publishing data on the Web.

the last being characterized by large amounts of un-structured as well as structured data in various forms, and ad-hoc approaches to data publishing, driven by Web-centric skills and technologies, without explicit support for semantic or spatial aspects. Since the data is often published over HTTP in various formats and with different structure, XML formats are often used to integrate different data structures from different resources into the Web.

Spatial data can be published on the Web in the JSON format through a RESTful API or through a Spatial Data Infrastructure (SDI) based on a Service Oriented Architecture (SOA). RDF data is also widely used to describe and link resources on the Web. In the spatial data domain, for example, OGC's GeoSPARQL is an RDF-based representation that is used to query spatial data.

1.2. Contributions

With the prevalence of sensor and actuator devices and the increase in location-based services, the use of spatial technologies is rapidly growing. The existing geospatial data, online maps combined with new forms of dynamic location-based data and services, create an opportunity for various new applications and services. However, to make spatial data more effectively available across different domains, a set of common practices are required.

The Spatial Data on the Web (SDW) Working Group⁷ has been committed to determining how spatial data can best be published, discovered, queried and integrated with other data on the Web. This paper is the result of considerable effort on identifying these best practices for publishing and integrating spatial data on the Web. These Spatial Data on the Web Best Practices were published as a W3C Note [3] and as an OGC Best Practice.

This companion paper was written by members of the Spatial Data on the Web Working Group. It summarizes the key requirements for publishing, retrieving and accessing spatial data on the Web to make it more interoperable, accessible, interpretable and understandable by humans and machines, accompanied by know-how and best practices addressing those requirements. The main contribution of the paper is its additional background information about the rationale underlying the selection of the best practices. In addition, the paper describes areas where best practices are still missing.

1.3. Paper organization

The rest of the paper is organized as follows: Section 2 explains a set of principles for setting out the scope of the problems that the best practices will address. Section 3 states the key requirements for publishing and sharing spatial data on the Web, presents the related best practices as identified by the working group, and discusses how the best practices address the described requirements. Section 4 discusses several gaps that still exist in current practice. Finally, Section 5 draws conclusions and discusses the future direction.

In the remainder of this paper, we refer to the Spatial Data on the Web Best Practices as "SDWBP".

2. Principles for describing best practices

As explained in the introduction, the aim of the work is to improve the discoverability, interoperability and accessibility of spatial data. The key principle follows from this: that through the adoption of the best practices identified in the SDWBP, the discoverability and linkability of spatial information published on the Web is improved.

⁷<http://www.w3.org/2015/spatial/>

A second principle concerns the intended audience of the spatial data in question. The aim is to deliver benefits to the broadest community of Web users possible—not to geospatial data experts only. The term ‘user’ signifies a data user: someone who uses data to build Web applications that provide information to end users—website visitors and app users—in some way. These data users are therefore among the intended audience of not only the spatial data, but also the Best Practice. The SDWBP provides value and guidance for application, website and tool builders to address the needs of the mass consumer market. Furthermore, the best practices should provide guidance for spatial data custodians. The best practices can offer a comprehensive set of guidelines for publishing spatial data on the Web.

A third principle is to have a broad focus. The first working draft of the best practices solicited several comments about a perceived “RDF bias”. While to develop spatial data following the 5-star Linked Data principles⁸ was one of the goals at the start, the solutions described in the best practices for discoverability and linkability should also be applicable to other spatial data on the Web. The best practices promote a linked data approach, but without asserting a strong association between linked data and RDF. Linked data requires only that the formats used to publish data support Web linking [28]. Furthermore, any ontologies developed within the working group should not be tightly coupled with upper ontologies (“compatible with” rather than “dependent upon”); this avoids data publishers having to commit to a given world view, as specified within a particular upper ontology, in order for them to use the best practices and any ontologies related to them.

A fourth principle follows from the term ‘best practice’ very directly: that its contents are taken from practice. The aim is not to reinvent or provide ‘best theories;’ in other words, the intention of the best practice is not to create new solutions where good solutions already exist or to invent solutions where they do not yet exist. The contents of the best practice should be made up of the best existing practices around publishing spatial data on the Web that can be found. Consequently, the aim is for each of the best practices in the document to be linked to at least one publicly available example(s) of a non-toy dataset that demonstrates the best practice.

Lastly, the best practices should comply with the principles of the W3C Best Practices for Publishing Linked Data [29] and the W3C Data on the Web Best Practices [30]. Where they do not, this will be identified and explained. The Data on the Web Best Practices (referred to as DWBP henceforth) form a natural counterpart to the work on spatial data on the Web. The best practices are aligned with DWBP in the following ways: a) by using the same best practice template, and b) by referring to the DWBP instead of repeating it. While the focus of the best practices is on spatial data, they may include recommendations on matters that are not exclusively related to spatial data on the Web, but are considered by the Working Group to be essential considerations in some use cases for publishing and consuming spatial data on the Web, and not covered in enough detail in DWBP or other documents.

3. The key requirements and best practices for publishing spatial data on the Web

The following sections discuss the main topics covered by SDWBP, and explains how they are addressed in the defined best practices. The topics are presented in a summarized fashion and are grouped thematically. Examples of datasets in which these best practices are implemented are not provided, since these are already present in the SDWBP. For the convenience of the reader, Table 2 outlines the best practices as they are stated in SDWBP, and indicates in which of the following sections they are discussed.

3.1. Geometries and spatial relationships

Tobler’s first law of geography states, “everything is related to everything else, but near things are more related than distant things” [31]. Like statistics, in addition to portraying information, spatial data provides a basis for reasoning, in this case based on location. Integration of data based on a location (i.e. relating things based on being located at or describing the same location) is often very useful. For this to be possible, either explicit geometries or topological relationships are necessary. Ideally, to enable their widest reuse, geometries should be described having in mind the geospatial, Linked Data and Web communities. This may not always be feasible, but the objective should at least be to describe geometries (also) for Web consumption.

⁸<http://5stardata.info>

Table 2
The Spatial Data on the Web Best Practices [3]—and where they are discussed in the paper

Best Practice 1	Use globally unique persistent HTTP URIs for spatial things	§3.3	Best Practice 2	Make your spatial data indexable by search engines	§3.9
Best Practice 3	Link resources together to create the Web of data	§3.3	Best Practice 4	Use spatial data encodings that match your target audience	§3.1, §3.6
Best Practice 5	Provide geometries on the Web in a usable way	§3.1, §3.10	Best Practice 6	Provide geometries at the right level of accuracy, precision, and size	§3.1, §3.5, §3.8
Best Practice 7	Choose coordinate reference systems to suit your user's applications	§3.2, §3.5, §3.7	Best Practice 8	State how coordinate values are encoded	§3.2
Best Practice 9	Describe relative positioning	§3.10	Best Practice 10	Use appropriate relation types to link spatial things	§3.1
Best Practice 11	Provide information on the changing nature of spatial things	§3.3, §3.6, §3.7	Best Practice 12	Expose spatial data through 'convenience APIs'	§3.3
Best Practice 13	Include spatial metadata in dataset metadata	§3.4	Best Practice 14	Describe the positional accuracy of spatial data	§3.5

One of the key best practices (namely, Best Practice 5) is therefore about providing geometries on the Web in a usable way. A single best way of publishing geometries was not identified; what is 'best' is in this case primarily related to the specific use case and tool support, which determine the geometry format to be used, the coordinate reference system (CRS) (more details are included in Section 3.2), as well as the level of accuracy, precision, size and dimensionality of geometry data. Note that these aspects are interrelated: for instance, the dimensionality of a geometry constrains the CRSs that can be used, as well as the geometry encodings.

Best Practice 5 identifies three scenarios in which geometries can be used: specific geospatial applications, linked data applications, and Web consumption. The practice also offers guidelines for choosing the right vocabularies from several available ones for describing geometry in each scenario. Currently, there are two geometry formats widely used in the geospatial and Web communities, respectively, GML [20] and GeoJSON [32]. GML provides the ability to express any type of geometry, in any CRS, and up to 3 dimensions (from points to solids) but is typically serialized in XML. GeoJSON supports only one coordinate reference system (CRS84—i.e., WGS 84 longitude / latitude), and geometries up to 2 dimensions (points, lines, surfaces) but it is serialized in JSON, which is often easier for browser-based Web applications to process. In the Linked Data community, several specific vocabularies for RDF-based representations of geometries are available, such as GeoSPARQL [27], W3C Basic

Geo [33], GeoRSS [34], or the ISA Core Location vocabulary [35]. Appendix A of SDWBP offers a comparison of the most common spatial ontologies.

Instead of catering for a single scenario, spatial data publishers should offer multiple geometric representations when possible, balancing the benefit of ease of use against the cost of the additional storage or additional processing if converting on-the-fly. This can be implemented using HTTP content-negotiation; however, this only works for media-type, character set, encoding and language. Consequently, it is not possible to select one representation that conforms to a given "profile" (e.g., data model, complexity level, CRS) from several that all share the same media-type.

Note that publishing geometries on the Web need not always be called for. Although spatial relationships can often be derived mathematically based on geometry, this can be computationally expensive. Topological relationships such as these can be asserted, thereby removing the need to do geometry-based calculations. Exposing such entity-level links to Web applications, user-agents and Web crawlers allows the relationships between resources to be found without the data user needing to download the entire dataset for local analysis.

3.2. Coordinate reference systems and projections

The key to reasoning and sharing spatial information is the establishment of a common coordinate reference frame in which to position the data. For Earth-based data, this may be done in spherical coordinate values

Fig. 2. Ellipsoid and spherical coordinates. (“Creative Commons Attribution 3.0 Australia” license from ICSM.gov.au—Intergovernmental Committee on Surveying and Mapping).

such as latitude, longitude and (optionally) elevation, or in a projected Cartesian coordinate space. The latter involves the flattening of a sphere in exchange for making it vastly easier to accurately measure area and distance. Regardless of the Coordinate Reference System (CRS) chosen, a distortion of the data will occur either in relative angles (positions), sizes (areas), or distances. Best Practice 7 identifies the World Geodetic System 1984 (WGS 84—EPSG:4326), which provides a good approximation at all locations on the Earth, as the most commonly used CRS for spatial data on the Web, but also explains when EPSG:4326 is not recommended, especially in use cases that require a higher level of accuracy than WGS 84 can offer.

In the Spatial Data on the Web Working Group, there was a lot of discussion on this topic. There are a lot of concerns with WGS 84 in the geospatial community. However, WGS 84 is dominant on the Web; it is used by most online mapping providers (e.g. Google, Bing, OpenStreetMap) and popular formats such as GeoJSON. Some explanation is required in order to understand the concerns of the geospatial community. All spatial coordinate frameworks begin with a mathematical model of the object being mapped. For the Earth, (which is not a true spheroid) an ellipsoid model, most fitting to the area being mapped, is commonly used (Figure 2). A Geodetic Datum is then placed on top of the reference ellipsoid to allow numerical expression of position. This may include a Vertical Datum, usually an approximation of sea level, to ref-

Upper left: Globular. Upper right: Orthographic. Lower left: Stereographic. Lower right: Mercator
What four commonly used projections do, as shown on a human head

Fig. 3. Projected CRS. (Public domain image from *Scientific American* circa 1923).

erence height and depth. WGS 84 (EPSG:6326) is an example of a Geodetic Datum based on the WGS 84 (EPSG:7030) Ellipsoid. It uses latitude and longitude coordinates (anchored to the equator and poles) to indicate position.

Geodetic CRSs are useful for collecting information within a common frame. The measurements they use are good for plotting directions but difficult to use when calculating area or distance. Latitude and longitude are angular measurements that do not convert easily to distance because their size in true units of measure (e.g., meters) varies according to location on the sphere. They become smaller as you near the poles.

In order to measure size and distance, and to display spatial information on a screen or paper, Projected Coordinate Reference Systems are used. A Projected CRS flattens a sphere to enable it to be portrayed and measured in 2 (or 3) dimensions. Figure 3 shows an example of the effects of this based on WGS 84 / Pseudo-Mercator (EPSG:3857), which is used commonly on the Web. As demonstrated by this figure, flattening a sphere introduces distortions. Imagine flattening the skin of an orange. You can preserve fairly accurately measurements for a portion of the skin but the rest will necessarily be distorted, either through stretching, compression or tears (see Figure 4). Projected CRSs are optimized to preserve distance, area or angular relations between spatial things for a chosen region. Distortion grows the further you move from that region. This is one reason why there are so many CRSs.

Fig. 4. The Web site <http://thetruesize.com/> is a good tool for comparing sizes of countries at different latitude and shows the distortion which results from flattening a sphere.

To summarize, WGS 84 is by far the most common CRS for spatial data on the Web. However, other CRSs exist for good reasons. Best Practice 7 therefore recommends that spatial data be published in CRSs to suit the potential user's applications. Spatial data on the Web should be published at least in WGS 84, and additionally in other CRSs if there are use cases demanding this.

The ability to unambiguously identify the CRS used is fundamental for the correct interpretation of spatial data. For instance, part of the information defined in a CRS concerns the order in which the geographic coordinates (i.e., latitude, longitude, etc.) are encoded, the units of measurement used for these coordinates, as well as the *datum* used. Mistaking the coordinate order (i.e., the *axis order*) is a very common error that results in plotting the data in a completely different location. Best Practice 8 specifically addresses this issue, by requiring that this information is made explicit, either by specifying the CRS used—use of EPSG codes, where they exist, is recommended⁹—, or by using data formats / vocabularies where this information is implicit (e.g., GeoJSON [36] and the W3C Basic Geo vocabu-

⁹EPSG is a register of CRSs maintained by the IOPG, an oil industry organization. The ESPG register is available online at: <http://www.epsg-registry.org/>.

lary [33] support only one CRS—respectively, CRS84 and WGS 84).

3.3. Spatial identifiers

Spatial things should be uniquely identified with persistent Uniform Resource Identifiers (URIs) in order for those using spatial data on the Web (Best Practice 1) to be able to definitively combine and reference these resources (Best Practice 3): they become part of the Web's information space; contributing to the Web of Data. URIs are a single global identification system used on the World Wide Web, similar to telephone numbers in a public switched telephone network. HTTP(S) URIs are a key technology to support Linked Data by offering a generic mechanism to identify entities ('Things') in the world and to allow referring to such entities by others. Anyone can assign identifiers to entities ('Things') in a namespace they own, i.e. a domain name within the Internet—e.g., hospitals, schools, roads, equipment, etc. 'Spatial things', such as the catchment area of a river, the boundaries of a building, a city or a continent, are examples of such 'Things' on the Web that need to be identified so that it is possible to refer to or make statements about this particular spatial thing.

Spatial things described or mentioned in a dataset on the Web should be identified using a globally unique URI so that a given spatial thing can be unambiguously identified in statements that refer to it. Good identifiers for data on the Web should be dereferenceable / resolvable, which makes it a good idea to use HTTP—or HTTPS—URIs as identifiers. Data publishers need to assign their subject spatial things HTTP URIs from an Internet domain name where they have authority over how the Web server responds. Typically, this means minting new HTTP URIs. Important aspects of this include authority, persistence, and the difference between information resources and the thing they give information about. The use of a particular Internet domain may reinforce the authority of the information served. HTTP can only serve information resources such as Web pages or JSON documents. Best Practice 1 contains no requirement to distinguish between the spatial thing and the page/document unless an application requires this.

HTTP URIs allow objects to link (through hyperlinking) with the existing objects on the Web (using their URIs), i.e., link to URIs available in community-maintained Web resources such as DBpe-

dia¹⁰ and GeoNames¹¹ or public government data, such as a national registers of addresses. Mapping and cadastral authorities maintain datasets that provide geospatial reference data. Re-using well-known identifiers is a good practice because it makes it easy to recognize that data from different sources are related. An example of a spatial thing URI is <http://sws.geonames.org/2172517/> which identifies the spatial thing *Canberra*, while when resolved in a browser returns information such as the name of the spatial thing, the centroid location, geographical boundaries, etc.

When exposing spatial data through standard SDIs (e.g., WFS services [10]), a certain user group is catered for, i.e., users with the expertise and tooling to use these services based on standards from the geospatial domain. To allow more users to benefit from the data it is important to expose the link to the Web representation of the features on the Web. Best Practice 12 identifies two approaches for doing this while keeping the SDI in place. One is to add an attribute to all spatial things, named `rdf_seealso`, which contains the URI of the Web representation of the spatial thing visible on the map. The Web representation is created by mapping the data in the SDI dynamically to crawlable resources on the Web using the R2RML standard [37] and Linked Data Publication tools. This approach leverages existing SDIs while enriching them with dereferenceable linked data representations of the spatial things. Exposing the data about a spatial thing as linked data makes sure that the attributes themselves will be URIs and thus unambiguous. The overhead of the extra attribute on existing SDIs is negligible and traditional clients will not be hindered by the extra attribute, but more advanced usage allows unlocking of a wealth of connections behind the traditional spatial data.

The other approach is to create a RESTful API as a wrapper, proxy or a shim layer around WFS services. It is worth mentioning that different spatial data services are often making their data accessible through REST API services. Content from the WFS service can be provided in this way as linked data, JSON or another Web-friendly format. There are examples of this approach of creating a convenience API that works dynamically on top of WFS such as the experimental *ldproxy*.¹² This is an attractive option for quickly

exposing spatial data from existing WFS services on the Web. The approach is to create an intermediate layer by introducing a proxy on top of the WFS (data service) and CSW (metadata service) [8] so the contained resources are made available. The proxy maps the data and metadata to schema.org [38] according to a provided mapping scheme; assigns URIs to all resources based on a pattern; makes each resource available in HTML, XML, JSON-LD [39], GML [20], GeoJSON [36], and RDF/XML (metadata only); and generates links to data in other datasets using SPARQL queries [40].

3.4. Discovery of spatial information

Cataloging of spatial information has always been difficult, whether the data is digital or not. A roadmap for Wellington NZ may be filed under NZ, Wellington, Transportation, tourism or a large number of other categories. Spatial data therefore has a greater requirement for metadata. Best Practice 13 recommends the inclusion of spatial metadata in dataset metadata. As a minimum, the spatial extent should be specified: the area of the world that the dataset describes. This information is necessary to enable spatial queries within catalog services such as those provided by SDIs and often suffices for initial discovery. However, further levels of description are needed for a user to be able to evaluate the suitability of a dataset for their intended application. This includes at least spatial coverage (continuity, resolution, properties), and representation (for example vector or grid coverage) as well as the coordinate reference system used.

In SDIs, the accepted standard for describing metadata is ISO 19115 [4, 5] or profiles thereof. To provide information about the spatial attributes of the dataset on the Web, DCAT [41] is recommended. An application profile of DCAT for geospatial data, GeoDCAT-AP [42], can be applied to more fully express spatial metadata. In addition, several spatial ontologies, already mentioned in Section 3.1, allow the description of spatial datasets.

3.5. Scale and quality

The quality of spatial data, as that of any other data, has a big impact on the quality of applications that use such data. This is intensified in the Web scenario, where data could be consumed with unforeseen purposes and data that are useful for a particular use case could come from plenty of sources.

¹⁰<http://dbpedia.org>

¹¹<http://geonames.org>

¹²<https://hub.docker.com/r/iide/ldproxy/>

Therefore, having quality information about spatial data on the Web significantly facilitates two main tasks to the consumer of such data. One of them is the selection of data, allowing to focus on data that satisfy the needs of a concrete use case. For example, spatial data accuracy is important when using them in the application of self-driving cars; guiding an autonomous vehicle to a precise parking spot near a facility that has a time-bounded service and then booking the charging spot requires accurate location data. Another is the reuse of data, i.e., understanding the behavior of data in order to adapt its processing (e.g., by considering data currentness, refresh rate or availability). A fundamental concept of spatial data is the scale of the representation of the spatial thing. This is important because combining data designed to be used at differing scales often produces misleading results. A scale is often represented as a ratio or shortened to the denominator value of a ratio. A 1:1,000 (or 1,000) scale map is referred to as larger than a 1:1,000,000 (or 1,000,000) scale map. Conversely, a million scale map is said to be a smaller scale than a thousand scale map. Data collected at a small scale is most often more generalized than data collected at a larger scale. Concepts related to scale are resolution (the smallest difference between adjacent positions that can be recorded), accuracy (the amount of uncertainty—how well a coordinate designates a position on the actual Earth’s surface) and precision (the reproducibility of a measurement to a certain number of decimal places in coordinate values).

When publishing spatial data on the Web, they should be supplemented with information about the precision and accuracy of such data. Such quality information should at least be available for humans. Best Practice 14 is concerned with positional accuracy and how spatial data should be accompanied with information about its accuracy. Best Practice 7 asserts a link between CRS and data quality, because the accuracy of spatial data depends for a large part on the CRS used, as was explained in a previous section. In order to support automatic machine-interpretation of quality information, such information should be published following the same principles as any other data published on the Web. The CRS of geometries on the Web should always be made known. For describing other quality aspects, the W3C Data Quality Vocabulary (DQV) [43], which allows specifying data quality information (such as precision and accuracy), could be used.

Even if the recommendations in Best Practice 14 focus on precision and accuracy, evidently the same ad-

vice can be followed for other relevant spatial quality information.

3.6. Thematic layering and spatial semantics

Spatial data is typically collected in layers. Although this sounds very map-oriented, these layers can be thought of as collections of instances of a class within a spatial and temporal frame. In other words, layers are usually organized semantically. Although SDWBP does not address layers directly, it does address spatial semantics. DWBP recommends the use of vocabularies to communicate the semantics, i.e., the meaning of data, and preferably standardized vocabularies. There are several vocabularies about spatial things available, such as the W3C Basic Geo vocabulary [33], GeoRSS [34], GeoSPARQL [27], the ISA Core Location Vocabulary [35] or schema.org [38]; overviews on their uses are provided in the literature [44–46]. Best Practice 4 identifies the main vocabularies in which spatial things can be described when the aim is data integration; however, it does not recommend one of them as the best. Currently there is no common practice in the sense of the same spatial vocabulary being used by most spatial data publishers. This depends on many factors; furthermore, describing spatial data multiple times using different vocabularies maximizes the potential for interoperability and lets the consumers choose which is the most useful. Appendix A of SDWBP offers guidance for selecting vocabularies by providing a table comparing the most commonly used ones.

The most important semantic statement to be made when publishing spatial data—or any data—is to specify the type of a resource. The W3C Basic Geo vocabulary has a class `SpatialThing` which has a very broad definition. This can be applicable (as a generic concept) to most of the common use-cases. Thematic semantics and general descriptions of spatial things and their properties should be provided as linked data. They should have URIs that return human- and machine-readable definitions when resolved.

3.7. Temporal dimension

It is important to consider the temporal dimension when publishing spatial data. The “where” component of the data is seldom independent of the “when” which may be implicit or explicit to the data. Hence, capturing the temporal component makes spatial data more useful to potential users, since it allows them to verify

whether the data suits their needs. For instance, tectonic movements over time can distort the coordinate values of spatial things.

This is included in Best Practice 7 on coordinate reference systems as valuable knowledge when dealing with spatial data for precision applications. Furthermore, it is recommended in Best Practice 13 to include metadata statements about the (most recent) publication date, the frequency of update and the time period for which the dataset is relevant (i.e., temporal extent).

Apart from the need for enhancing spatial data with their temporal context, the temporal dimension also affects the very nature of spatial things, since both spatial things and their attributes can change over time; this is covered in Best Practice 11. When dealing with changes to a spatial thing, its lifecycle should be taken into account; in particular, how much change is acceptable before a spatial thing can no longer be considered as the same resource (and requires defining a new resource with a new identifier). Creating a new resource will depend on whether domain experts think the fundamental nature of the spatial thing has changed, taking into account its lifecycle (e.g., a historic building replaced by another that has been built on top of it). In this case, the temporal dimension of spatial data can be expressed by providing a series of immutable snapshots that describe the spatial thing at various points in its lifecycle, each snapshot having a persistent URI.

In those cases when the spatial thing itself does not change over time but its attributes do, the description of the spatial thing should be updated to reflect these changes, and each change published as a snapshot. In contrast, when a spatial thing has a small number of attributes that are frequently updated (e.g., the GPS-position of a runner or the water level from a stream gauge), the time-series of data values within such attributes of the spatial thing should be represented. This is relevant in relation to recent advances in embedding smart sensors and actuators in physical objects and machines such as vehicles, buildings, and home appliances, which has led to the publishing of large volumes of data that are spatio-temporal [47].

3.8. Size of spatial datasets

Spatial data tends to be very large. This can pose difficulties when sharing or consuming spatial data over the Web—particularly in low bandwidth or high latency situations. Accurate (polygon) geometries tend to contain a high number of coordinates. Especially when querying collections of spatial things with ge-

ometries over the Web, this results in very large response payloads wasting bandwidth and causing slow response times, while for some very common use cases, like simply displaying some things on a Web map, high accuracy is not required. The primary basis for simplification is scale. For example, when searching for a Starbucks on a city scale, an accuracy of 3 meters is acceptable, but when providing street-level directions to a shop for a self-driving car it is not.

For those use cases that do not require high accuracy, common ways of dealing with reducing the size of spatial data include degrading precision by reducing the number of decimals, and simplifying geometries using a simplification algorithm such as Ramer-Douglas-Peucker [48, 49] or Visvalingam-Whyatt [50]. These methods result in lower accuracy and precision.

Big spatial data is often not vector data, i.e., a representation of spatial things using points, lines, and polygons [16], but coverage data, i.e., gridded data: a data structure that maps points in space and time to property values [51]. For example, an aerial photograph can be thought of as a coverage that maps positions on the ground to colors. A river gauge maps points in time to flow values. A weather forecast maps points in space and time to values of temperature, wind speed, humidity and so forth. For coverage data, other methods are required to manage size.

DWBP recommends to provide 1) bulk download and 2) subsets of data. Providing bulk-download or streaming access to data is useful in any case and is relatively inexpensive to support as it relies on standard capabilities of Web servers for datasets that may be published as downloadable files stored on a server. Subsets, i.e., extracts or “tiles”, can be provided by having identifiers for conveniently sized subsets of large datasets that Web applications can work with [52]. Actually, breaking up a large coverage into pre-defined lumps that you can access via HTTP GET requests is a very simple API. A second way of supporting extracts, more appropriate for frequently changing datasets, is by supplying filtering options that return appropriately sized subsets of the specific dataset.

3.9. Crawlability

Search engines use crawlers to update their search index with new information that has been found on the Web. Traditional crawlers consist of two components: URL extractors, i.e., HTML crawlers that extract links

from HTML pages in order to find additional sources to crawl, and indexers, i.e., different types of indexes, typically using the occurrence of text on a Web page, that are maintained by search engines.

A major issue with crawlers identified in the early 2000's by Bergman [53] was the inaccessibility of the so-called "deep Web": information that was hidden to traditional crawlers as it is only accessible through services (e.g., forms) that require user input. Several solutions have since been introduced to access and index information on the "deep Web" [54, 55]. Spatial Data services (e.g. OGC Web services and/or other APIs) typically make information available only after user input has been provided, leading to a similar problem, i.e., a Deep Spatial Web. Further, these services are built to be accessed and searched by domain-specific applications rather than general Web services and/or search engine crawlers. For example, in the OGC architecture, catalog services are intended to be used for searching spatial assets by using specific client tools, and are not optimized for indexing and discovery by general purpose search engines. However, a typical Web user does not know that these catalogs exist and is accustomed to using general purpose search engines for finding information on the Web. Therefore, making sure that spatial data is indexable by Web search engines is an important approach for making spatial data discoverable by users directly. The addition of structured data to Web services that are otherwise not accessible to search engines increases the visibility of a service or dataset in major search engines [56]. Schema.org [38], a single schema across a wide range of topics that includes people, places, events, products, offers, etc., and widely supported by Bing, Google, Yahoo! and Yandex, is the predominant way of marking up content and services on the Web with structured data to improve the presentation of the result in a search engine [56]. To verify if schema.org markup on a Web page is recognized by Web agents, Google's Structured Data testing tool¹³ can be used. Experimental work such as Geonovum's Spatial Data on the Web Testbed¹⁴ describes ways to make spatial data indexable by publishing an HTML Web page for the spatial dataset and each spatial thing that it contains; by using structured data, schema.org and links, as well as publishing XML sitemaps containing links to all data resources for spatial data services. Another example is

the Dutch geoportal PDOK.nl which extensively publishes dataset metadata, for example, the national roads dataset,¹⁵ resulting in better accessibility through common search engines. Several examples of spatial things published in this way are provided in Best Practice 2.

Currently, spatial information, even when published in accordance with these guidelines, is not widely exploited by search engines. However, by increasing the volume of spatial information presented to search engines, and the consistency with which it is provided, it is expected that search engines will begin offering spatial search functions. Evidence is already seen of this in the form of contextual search, such as prioritization of search results from nearby entities. In addition, search engines are beginning to offer more structured, custom searches that return only results that include certain schema.org types such as *Dataset*, *Place* or *City*.

3.10. Other aspects of spatial data

Spatial things may have 0 to 3 dimensions (points, lines, areas, 3D), and it may be difficult to combine similar things if the dimensions in which they are represented differ. Although SDWBP does not address this at length, Best Practice 5 does recommend describing the number of dimensions in metadata and notes that one of the selection criteria for choosing a geometry format on the Web is the dimensionality of the data.

In common language, and for spatial things inherently related to mobile things, it may be convenient to describe positions of spatial things *relative* to other spatial things. Just as for other spatial things, we would like such descriptions to be both human- and machine-interpretable. We advise, in Best Practice 9, that positions or geometries of the target reference things should be retrievable via link relations, in accord with general principles for linkability in Section 2. The geocentric use case (i.e., position relative to the Earth) is generally addressed throughout SDWBP and does not require an explicit link relation to the Earth. As the active contributors to SDWBP were primarily interested in the geocentric case, SDWBP does not explore relative positioning in depth. We find that practices in this area vary widely in the details of implementation and so we are able to offer only broad advice and examples. Spatial relationships as described for Best Practice 10

¹³<https://search.google.com/structured-data/testing-tool>

¹⁴<https://github.com/geo4web-testbed/general>

¹⁵<https://data.pdok.nl/datasets/nationaal-wegenbestand-wegen>

(Section 3.1) may be useful, but we find no evidence for suitable common vocabularies for many needs. For some spatial data, the *symbology* associated with the data is of high importance because it communicates meaning. However, as rendering maps is explicitly out of scope, symbology is not addressed in SDWBP.

4. Gaps in current practice

The best practices described in brief in Section 3, and in full in the SDWBP document, are compiled based on evidence of real-world application. This is in line with the fourth principle described in Section 2 of this paper. However, there are several issues that inhibit the use or interoperability of spatial data on the Web, for which no evidence of real-world applied solutions was found. These issues are denoted “gaps in current practice” and described as such in this section. An issue is considered a gap when no evidence of real-world applied solutions in production environments was found. The term ‘production environment’ signifies a case where spatial data has been delivered on the Web with the intention of being used by end users and with a quality level expected from such data. In contrast, a “testing environment” is published with the intent of being tested so that bugs can be discovered and fixed and an experimental publication of spatial data on the Web is published with the intent of, for example, exploring possibilities, learning about the technology, or other goals besides publishing with the intent of serious use. In the case of gaps, there might be emerging practice, i.e., a solution that has been theorized for a certain issue and has possibly been experimented on in testing/beta settings, but not in production environments. Gaps and emerging practices in the area of publishing spatial data on the Web are discussed in this section.

4.1. Representing geometry on the Web

Location information can be an important ‘hook’ for finding information and for integrating different datasets. There are different ways of describing the location of spatial things: referencing the name of a well-known named place, describing a location in relation to another location, or providing the location’s coordinates as a geometry. The latter allows the integration of data based on location using spatial reasoning, even when explicit links between things are not available, as well as, of course, showing spatial things on

a map. Although there are ways to represent geometry in Web data, there are still some gaps in current practice related to selecting a serialization format, selecting an embedded, file-based or Linked Data approach, and making geometries available in different CRSs.

There are several aspects to representing geometries on the Web. First, there is the question of different serialization formats to choose from. In general, the formats are the same as for publishing any other data on the Web: XML, JSON, CSV, RDF, etc. How to select the most appropriate serialization is described in general terms in DWBP. As described in Section 3.1, a single best way of publishing geometries was not identified in SDWBP. The currently most widely used geometry formats, namely, GML [20] and GeoJSON [36], do not address all requirements. Therefore, in order to facilitate the use of geometry data on the Web, SDWBP recommends that GML-encoded geometries are made available also in GeoJSON, by applying not only the required coordinate reference system transformation, but, if needed, by simplifying the original geometry (e.g., by transforming a 3D geometry in a 2D one).

A second, related aspect concerns the existing options for publishing geometries on the Web—e.g., in self-contained files such as GML or GeoJSON, or rather to embed geometries as structured data markup in HTML, or in an RDF-based way, i.e., as Linked Data. Choosing between these approaches—or not choosing but rather offering a combination of these—depends largely on the intended audience. As explained in Section 3.9, dealing with crawlability, the advantage is that HTML with embedded data is indexed by search engine crawlers. However, the options for embedding geometry in HTML are limited. Typically this is done using schema.org [38] as Microdata [57], RDFa [58], or JSON-LD [39], but for specifying only 0D-2D geometries (points, lines, surfaces)—e.g., the centroid and/or 2D bounding box. An additional issue is that search engines index only a subset of the terms defined in schema.org, and those concerning geometries are not fully supported. RDF-based publication of geometry data is the most advanced option, but the audience for this is smaller than the others. GeoSPARQL [27] offers a vocabulary that allows serialization of geometries as GML or WKT (“Well-Known Text”) [59], whereas the ISA Core Location vocabulary [35] supports also GeoJSON, but the lack of best practices on the consistent use of the existing spatial data vocabularies prevents interoperability (see Section 4.2).

Third, there is a question of how to make geometries available in different CRSs. Section 3.2 explains the existence of many CRSs and why spatial data should be published in CRSs that are most common to potential users. It follows that, on the one hand, the CRS should be specified for geometries published on the Web and, on the other hand, users should be able to find out which CRSs are available and to get geometries using the CRS of their choice.

Sometimes the CRS used is clear from the representation. In other cases the CRS needs to be specified either on the dataset level or the instance level. How this is done differs for each serialization. For example, in GeoSPARQL this is added as a prefix of the WKT literal while in GML an attribute `srsName` can be specified on geometry elements. In an OGC WFS [10] request, users can specify the CRS they wish to use by specifying the `srsName` parameter. In GeoSPARQL the `getSRID` function returns the spatial reference system of a geometry, thus making it possible to request a specific CRS at a (Geo)SPARQL endpoint. However, these options require the user to be proficient in either geospatial Web services or Linked Data.

A best practice for returning geometries in a specific requested CRS has not yet emerged. Many options can be found in current practice, including creating CRS-specific geometry properties (for example, the Dutch Land Registry does this), and supporting an option for requesting a specific CRS in a convenience API; but one best practice cannot yet be identified. Another option worth exploring might be the use of content negotiation, i.e., negotiate CRS as part of the content format for the geometry, as has also been proposed for encoding format. For example, this could be done with an extra media type parameter (e.g., `application/ttl; geomLiteral="WKT"; crs="CRS84"`) or by adding specific request and response headers for negotiating CRS to the HTTP protocol. A contribution to address this issue might be provided by the W3C Dataset Exchange Working Group (DXWG), which is due to deliver a specification on profile-based content negotiation.¹⁶ However, providing different CRS might be too complicated to handle in the HTTP protocol. For example, multidimensional datasets will in general use multiple CRSs (e.g., horizontal, vertical and temporal, maybe more), and conversion between

CRSs will, in general, introduce errors, so data in one CRS are not exactly the same as data in another CRS. Furthermore, CRS is just one of a number of parameters that may characterize a particular geometric representation of a spatial thing, including the type of geometry, its relationship to the Thing, method of interpolation, scale or resolution. However, offering a choice between all these parameters of data objects such as geometries might be an overloading of HTTP content negotiation protocols. It might, therefore, be more appropriate to handle this in the application layer.

4.2. A spatial data vocabulary

Although a large amount of geospatial data has been published on the Web, so far there are few authoritative datasets containing geometrical descriptions of spatial things available in Web-friendly formats. Their number is growing (e.g., at the time of writing there are three authoritative spatial datasets publicly available as linked data in the Netherlands containing topographic¹⁷, cadastral¹⁸, and address¹⁹ data), but currently there is no common practice in the sense of the same spatial vocabulary—or the same combination of spatial vocabularies—being used by most spatial data publishers. The consequence is the lack of a baseline during the mapping process for application developers trying to consume specific incoming data. Identifying administrative units, points of interest or postal addresses with URIs could be beneficial not only for georeferencing other datasets, but also for interlinking datasets georeferenced by the direct and indirect location information.

Direct georeferencing of data implies representing coordinates or geometries and associating them to a CRS. This requires vocabularies for geometries and able to specify which CRS is used. Further, indirect georeferencing of data implies associating them to other data on named places. Preferably, these data on named places should be also georeferenced by coordinates in order to serve as the basis for data linking between indirectly and directly georeferenced datasets. Moreover, vocabularies developed for representing specific sets of geospatial data on the Web should reuse as much as possible existing ones. This is the case for the vocabularies developed by IGN France for geome-

¹⁶For a description of this deliverable, see the DXWG Charter: <https://www.w3.org/2017/dxwg/charter>.

¹⁷<https://brt.basisregistraties.overheid.nl>

¹⁸<https://brk.basisregistraties.overheid.nl>

¹⁹<https://bag.basisregistraties.overheid.nl>

tries,²⁰ topographic entities,²¹ and CRS.²² These vocabularies contain alignments with existing vocabularies, e.g., the class `geom:Geometry` is a subclass of both `sf:Geometry` (OGC Simple Features vocabulary) which is a subclass of the `Geometry` class of the `GeoSPARQL` vocabulary; and `ngeo:Geometry` (Neogeo vocabulary). Furthermore, the topographic entity class from the IGN France vocabulary is declared equivalent to the `Feature` class from the `Geonames` vocabulary.

In W3C Basic Geo [33], it is assumed that the CRS used is WGS 84. However, publishers might have data in a different, local CRS. Thus, there is a need for a more generic class for, for example, a point geometry with the benefit of choosing the CRS of the underlying data [60]. Existing vocabularies, as `GeoSPARQL` [27] and `ISA Core Location` [35], support this feature, but there are currently no best practices for their consistent use, thus hindering interoperability.

Vocabularies like the one by IGN France are created because, currently, the existing vocabularies do not cover all requirements and no guidance is available on their consistent and complementary use. SDWBP partially addresses the latter issue, by providing examples on how to model spatial information with widely used vocabularies. However, solving the existing gaps would require the definition of new terms in existing or new vocabularies, which was not in scope with the work of the Spatial Data on the Web Working Group. A possible way forward is an update for `GeoSPARQL`. This would provide an agreed spatial ontology, i.e., a bridge or common ground between geographical and non-geographical spatial data and between W3C and OGC standards, conformant to the ISO 19107 [17] abstract model and based on existing vocabularies such as `GeoSPARQL` [27], the W3C Basic Geo Vocabulary [33], `GeoRSS` [34], `NeoGeo` or the `ISA Core Location` vocabulary [35]. The updated `GeoSPARQL` vocabulary would define basic semantics for the concept of a reference system for spatial coordinates, a basic datatype, or basic datatypes for geometry, how geometry and real world objects are related and how different versions of geometries for a single real world object can be distinguished. For example, it makes sense to publish different geometric representations of a spatial object that can be used for different purposes. The same object could be modeled

as a point, a 2D or a 3D polygon. The polygons could have different versions with different resolutions (generalization levels). And all those different geometries could be published with different coordinate reference systems. Thus, the vocabulary would provide a foundation for harmonization of the many different geometry encodings that exist today. An alternative approach is to establish best practices for the consistent use of the most popular spatial vocabularies. An example are the guidelines for the RDF representation of INSPIRE data developed in the framework of the EU ISA Programme [61] by following the SDWBP recommendations. In such a scenario, the definition of new classes and properties would be limited to cover the gaps in the existing vocabularies.

4.3. Spatial aspects for metadata

Even if all spatial data should become discoverable directly through search engines, data portals would still remain important hubs for data discovery—for example, because the metadata records registered there can be made crawlable. But, in addition, different data portals can harvest each others' information provided there is consistency in the types and meaning of included information, even if structures and technologies vary. Discovery of spatial data is improved in the Netherlands, for example, because the national general data portal²³ harvests the spatial data portal and thus all spatial datasets are registered in the general data portal as well.

In the eGovernment sector, DCAT [41] is a standard for dataset metadata publication, and harvesting this metadata is implemented by eGovernment data portals. Because DCAT is lacking in possibilities for describing some specific characteristics of spatial datasets, an application profile for spatial data, `GeoDCAT-AP` [42, 62], has been developed in the framework of the ISA Programme of the European Union,²⁴ with the primary purpose of enabling the sharing of spatial metadata across domains and catalog platforms. To achieve this, `GeoDCAT-AP` defines RDF bindings covering the core profile of ISO 19115:2003 [4] and the INSPIRE [63] metadata schema, enabling the harmonized RDF representation of existing spatial metadata. The focus was on the most used metadata elements, whereas additional mappings—as well as the alignment with the latest version of ISO 19115 [5]—could be defined in fu-

²⁰<http://data.ign.fr/def/geometrie/20160628.en.htm>

²¹<http://data.ign.fr/def/topo/20140416.en.htm>

²²<http://data.ign.fr/def/ignf/20160628.en.htm>

²³<https://data.overheid.nl>

²⁴<https://ec.europa.eu/isa2/>

ture versions of the specification, based on users' and implementation feedback.

One of the outcomes of the development of GeoDCAT-AP was the identification of gaps in existing RDF vocabularies for representing some spatial information [64]—such as coordinate reference systems and spatial resolution (see also Section 4.2 on this topic). But it also highlighted a key issue for spatial data, in that the use of global and persistent identifiers is far from being a common practice. Apart from making it difficult to implement a Linked Data-based approach, this situation has negative effects on the geospatial infrastructure itself. E.g., it makes it impossible to unambiguously identify a spatial thing or a dataset over time, and it prevents an effective implementation of incremental metadata harvesting in a federated infrastructure (such as the INSPIRE one).

Notably, recent activities are contributing to fill at least part of these gaps. For instance, DQV [43] provides a solution for modeling precision and accuracy, as mentioned in Section 3.5. Moreover, the use cases collected by the W3C Dataset Exchange Working Group [65] cover also geospatial requirements, which might be addressed by the work of this group in the revision to DCAT [41]. However, the consistent use of global and persistent identifiers in the geospatial domain is an issue that, far from being merely technical, affects the data management workflow, and therefore needs to be addressed also at the organizational level.

4.4. Describing dataset structure and service behaviors

Datasets may be arbitrarily large and complex, and may be exposed via services to expose useful resources, rather than a “download” scenario. Data gathered using automated sensors, in particular, may be impossible to download in its entirety due to its dynamic nature and potential volumes. It is, therefore, necessary in these cases to be able to adequately describe the structure of such data and how services interact to expose subsets of it—even individual records in a Linked Data context. Such datasets are common in the information processing world, and commonly organized in “hypercubes”—where “data dimensions” are used to locate values holding results. A standard based on this dimensional model of data is the RDF Data Cube vocabulary (QB) [66]. It has been used to publish sensor data; for example to publish a homogenized daily temperature dataset for Australia over the last 100 years [67]. However, QB is lacking in possibilities

for describing spatio-temporal aspects of data, which are very important for observations. One of the work items in the Spatial Data on the Web Working Group was, therefore, an extension to the existing QB vocabulary to support specification of key metadata required to interpret spatio-temporal data, called QB4ST [68].

QB4ST is an extension to QB to provide mechanisms for defining spatio-temporal aspects of dimension and measure descriptions. It is intended to enable the development of semantic descriptions of specific spatio-temporal data elements by appropriate communities of interest, rather than to enumerate a static list of such definitions. It provides a minimal ontology of spatio-temporal properties and defines abstract classes for data cube components (i.e., dimensions and measures) that use these, to allow classification and discovery of specialized component definitions using general terms. QB4ST is designed to support the publication of consistently described re-usable and comparable definitions of spatial and temporal data elements by appropriate communities of practice. One obvious such case is the use of GPS coordinates described as decimal latitude and longitude measures. Another example is the intended publication of a register of Discrete Global Grid Systems (DGGS) by the OGC DGGS Working Group. QB4ST is intended to support publication of descriptions of such data using a common set of attributes that can be attached to a property description (extending the available QB mechanisms for attributes of observations). The Spatial Data on the Web Working Group has demonstrated the use of QB and QB4ST to serve satellite imagery through DGGS and a virtualized triple store [52].

4.5. Versioning of spatial data

Future Internet technologies will aim more and more to capture, make sense of and represent not only static but dynamic content and up-to-date information. Through Internet technology, connections between devices and people will be realized through the exchange of large volumes of multimedia and data content. When the communication latency becomes lower, and the capacity of the communication gets higher with fifth generation (5G) mobile communications systems, the Internet will be an even more prominent platform to control real and virtual objects in different parts of our lives, such as healthcare, education, manufacturing, smart grids, and many more [69]. The Internet of Things [70] and Tactile Internet [71] are some of the technologies that aim to facilitate the interaction be-

tween people and devices, observe near real-time phenomena and actuate devices or robots. The Tactile Internet is focused on speeding up this interaction process and reducing the latency in communication systems. Such high-speed communication will bring new challenges for intelligent systems. There is a big gap in the lightweight and semantically rich representation of versioning and temporal aspects of spatial data content. There have been a few attempts to represent changing and moving spatial objects, such as OGC's TimeSeriesML [72] and Moving Features [73]. However, although these ontologies provide a reasonably good semantic coverage, there is still a need for the development of lightweight and semantically rich representations to conduct enhanced (near) real-time operations. Having heavy semantic expressivity in ontologies can cause a burden on reasoning engines and can slow down the processing time for machines. For instance, it is important to identify and represent the direction and coverage area of a surveillance camera or the orientations and positions of objects or people in the observed environment. In the future, a broken car will be fixed by a robot, or surgeries will be carried out by multiple robots using Tactile Internet [74]. It can also be envisioned that people, who have difficulty in walking, will not need to use a walking stick but merely a strap of an exoskeleton. These must be controlled by wireless systems to monitor the coverage, direction, and identify the objects and people around them including their shapes.

To conduct such activities, a better representation is needed for not only spatial information but also temporal and geometrical aspects of objects. Observed objects can change their size during the actuation process. For instance, a group of surgical robots will need to know about the shape of an organ and the changes regarding its size, geometrical shape, and orientation during surgery and exchange this information among themselves and with doctors to conduct an operation with high precision and low latency. A self-driving wheelchair or a self-driving car will be able to communicate with other sensor objects regarding the surrounding environment and direction to avoid obstacles. This will prevent possible accidents and harm caused by machines, such as not falling down stairs or running into objects with high speed or force. In all these scenarios, lightweight representation and exchange of temporal-spatial knowledge are essential to understand and react fast enough to prevent disasters or to control the movement of devices. Having the means to represent the semantics of activities and phenomena at such

high granularity and lightweight format will play a pivotal role in the development of future Internet technologies. Moreover, it will allow machines to instantly exchange information including spatial and geometrical knowledge and carry out their tasks with high precision. The Spatial Data on the Web Interest Group²⁵, the successor of the Working Group, will address the topic of representing moving objects on the Web.

5. Conclusions

Spatial data has become ubiquitous with the explosive growth in positioning technologies attached to mobile vehicles, portable devices, and autonomous systems. It has proven to be fundamentally useful for countless things, ranging from everyday tasks like finding the best route to a location to solving the biggest global challenges like climate change adaptation. However, spatial data dissemination is heterogeneous and although the Web is commonly used as a publication medium, the discovery, retrieval, and interpretation of spatial data on the Web is still problematic.

SDWBP describes how Web principles can be applied to the world of complex spatial data to solve this problem. Good practices can be observed in current practice and have been collected into the Best Practices based on a set of principles and an examination of practice. In some cases, a best practice has not yet emerged. There are still questions related to representing geometry on the Web, with regard to recommendable serialization forms and formats, and the use of coordinate reference systems. A Web-friendly way of publishing spatial metadata has not yet been described in full, especially with regards to the relevant subset of spatial metadata standards. A standardized ontology for spatial things that covers all the main requirements for publishing spatial linked data is not yet available, and best practices on the consistent use of the existing spatial vocabularies are yet to be established. Finally, there are new approaches emerging such as QB4ST [68], an extension to the RDF Data Cube to provide mechanisms for defining spatio-temporal aspects of dimension and measure descriptions.

Notwithstanding these gaps and emerging solutions, a useful set of actionable best practices for publishing spatial data on the Web has been described. Following these guidelines will enable data users, Web applica-

²⁵<https://www.w3.org/2017/sdwig/charter.html>

tions and services to discover, interpret and use spatial data in large and distributed Web systems.

Acknowledgements

The authors express their sincere gratitude to all the participants in the W3C/OGC Spatial Data on the Web Working Group, who contributed to the work discussed in this paper. The authors acknowledge partial support from NSF (award number 1540849).

References

- [1] K. Taylor and E. Parsons, Where Is Everywhere: Bringing Location to the Web, *IEEE Internet Computing* **19**(2) (2015), 83–87. doi:10.1109/MIC.2015.50. <https://doi.org/10.1109/MIC.2015.50>.
- [2] I. Masser, All shapes and sizes: the first generation of national spatial data infrastructures, *International Journal of Geographical Information Science* **13**(1) (1999), 67–84. doi:10.1080/136588199241463. <https://doi.org/10.1080/136588199241463>.
- [3] J. Tandy, L. van den Brink and P. Barnaghi, Spatial Data on the Web Best Practices, W3C Working Group Note, W3C, 2017. <https://www.w3.org/TR/2017/NOTE-sdw-bp-20170928/>.
- [4] ISO/TC 211, ISO 19115:2003 Geographic information – Metadata, ISO Standard, International Organization for Standardization (ISO), 2003. <https://www.iso.org/standard/26020.html>.
- [5] ISO/TC 211, ISO 19115-1:2014 Geographic information – Metadata – Part 1: Fundamentals, ISO Standard, International Organization for Standardization (ISO), 2014. <https://www.iso.org/standard/53798.html>.
- [6] ISO/TC 211, ISO 19119:2016 Geographic information – Services, ISO Standard, International Organization for Standardization (ISO), 2005. <https://www.iso.org/standard/39890.html>.
- [7] ISO/TC 211, ISO 19119:2016 Geographic information – Services, ISO Standard, International Organization for Standardization (ISO), 2016. <https://www.iso.org/standard/59221.html>.
- [8] D. Nebert, U. Voges and L. Bigagli, OGC® Catalogue Services 3.0 - General Model, OGC® Implementation Standard, OGC, 2016. <http://www.opengis.net/doc/IS/cat/3.0>.
- [9] J. de la Beaujardiere, OpenGIS Web Map Server Implementation Specification, OpenGIS® Implementation Standard, OGC, 2006. <http://www.opengis.net/doc/IS/wms/1.3>.
- [10] P.A. Vretanos, OpenGIS Web Feature Service 2.0 Interface Standard, OpenGIS® Implementation Standard, OGC, 2010. <http://www.opengis.net/doc/IS/wfs/2.0>.
- [11] P. Baumann, OGS® Web WCS 2.0 Interface Standard - Core, version 2.0.1, OGS® Implementation Standard, OGC, 2012. <http://www.opengis.net/doc/IS/wcs-core-2.0.1>.
- [12] ISO/TC 211, ISO 19117:2012 Geographic information – Portrayal, ISO Standard, International Organization for Standardization (ISO), 2012. <https://www.iso.org/standard/46226.html>.
- [13] M. Lupp, Styled Layer Descriptor profile of the Web Map Service Implementation Specification, OGC® Implementation Standard, OGC, 2007. <http://www.opengis.net/doc/IS/sld/1.1>.
- [14] M. Müller, Symbology Encoding Implementation Specification, OpenGIS® Implementation Standard, OGC, 2006. <http://www.opengis.net/doc/IS/se/1.1>.
- [15] D. Burggraf, OGC KML 2.3, OGC® Implementation Standard, OGC, 2015. <http://www.opengis.net/doc/IS/kml/2.3>.
- [16] ISO/TC 211, ISO 19103:2015 Geographic information – Conceptual schema language, ISO Standard, International Organization for Standardization (ISO), 2015. <https://www.iso.org/standard/56734.html>.
- [17] ISO/TC 211, ISO 19107:2003 Geographic information – Spatial schema, ISO Standard, International Organization for Standardization (ISO), 2003. <https://www.iso.org/standard/26012.html>.
- [18] ISO/TC 211, ISO 19109:2015 Geographic information – Rules for application schema, ISO Standard, International Organization for Standardization (ISO), 2015. <https://www.iso.org/standard/59193.html>.
- [19] ISO/TC 211, ISO 19110:2005 Geographic information – Methodology for feature cataloguing, ISO Standard, International Organization for Standardization (ISO), 2005. <https://www.iso.org/standard/39965.html>.
- [20] C. Portele, OpenGIS® Geography Markup Language (GML) Encoding Standard. Version: 3.2.1, OpenGIS® Standard, OGC, 2007. <http://www.opengis.net/doc/IS/gml/3.2.1>.
- [21] ISO/TC 211, ISO 19111:2005 Geographic information – Spatial referencing by coordinates, ISO Standard, International Organization for Standardization (ISO), 2005. <https://www.iso.org/standard/26016.html>.
- [22] ISO/TC 211, ISO 19111:2007 Geographic information – Spatial referencing by coordinates, ISO Standard, International Organization for Standardization (ISO), 2007. <https://www.iso.org/standard/41126.html>.
- [23] F. Knibbe and A. Llaves, Spatial Data on the Web Use Cases & Requirements, W3C Working Group Note, W3C, 2016. <https://www.w3.org/TR/2016/NOTE-sdw-ucr-20161025/>.
- [24] Y. Fathy, P. Barnaghi and R. Tafazolli, Large-Scale Indexing, Discovery and Ranking for the Internet of Things (IoT), *ACM Computing Surveys* (2017), (in press). <http://epubs.surrey.ac.uk/id/eprint/844698>.
- [25] J. Noguera-Iso, F.J. Zarazaga-Soria, R. Béjar, P.J. Álvarez and P.R. Muro-Medrano, OGC Catalog Services: a key element for the development of Spatial Data Infrastructures, *Computers & Geosciences* **31**(2) (2005), 199–209. Geospatial Research in Europe: AGILE 2003. doi:10.1016/j.cageo.2004.05.015. <https://doi.org/10.1016/j.cageo.2004.05.015>.
- [26] R. Cyganiak, D. Wood and M. Lanthaler, RDF 1.1 Concepts and Abstract Syntax, W3C Recommendation, W3C, 2014. <http://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/>.
- [27] M. Perry and J. Herring, GeoSPARQL - A Geographic Query Language for RDF Data. Version 1.0, OGC® Implementation Standard, OGC, 2010. <http://www.opengis.net/doc/IS/geosparql/1.0>.
- [28] I. Jacobs and N. Walsh, Architecture of the World Wide Web, Volume One, W3C Recommendation, W3C, 2004. <http://www.w3.org/TR/2004/REC-webarch-20041215/>.
- [29] B. Hyland, G. Atemez and B. Villazón-Terrazas, Best Practices for Publishing Linked Data, W3C Working Group Note, W3C, 2014. <http://www.w3.org/TR/2014/NOTE-ld-bp-20140109/>.

- [30] B. Farias Lóscio, C. Burle and N. Calegari, Data on the Web Best Practices, W3C Recommendation, W3C, 2017. <https://www.w3.org/TR/2017/REC-dwbp-20170131/>.
- [31] W.R. Tobler, A Computer Movie Simulating Urban Growth in the Detroit Region, *Economic Geography* **46**(sup1) (1970), 234–240. doi:10.2307/143141. <https://doi.org/10.2307/143141>.
- [32] W. Li, M. Song, B. Zhou, K. Cao and S. Gao, Performance improvement techniques for geospatial web services in a cyberinfrastructure environment - A case study with a disaster management portal, *Computers, Environment and Urban Systems* **54** (2015), 314–325. doi:10.1016/j.compenvurbsys.2015.04.003. <https://doi.org/10.1016/j.compenvurbsys.2015.04.003>.
- [33] D. Brickley, Basic Geo (WGS84 lat/long) Vocabulary, W3C Incubator Group Report, W3C, 2006. <https://www.w3.org/2003/01/geo/>.
- [34] J. Lieberman, R. Singh and C. Goad, W3C Geospatial Vocabulary, W3C Incubator Group Report, W3C, 2007. <https://www.w3.org/2005/Incubator/geo/XGR-geo/>.
- [35] A. Perego and M. Lutz, ISA Programme Location Core Vocabulary, Namespace Document, European Commission, 2015. <https://www.w3.org/ns/locn>.
- [36] H. Butler, M. Daly, A. Doyle, S. Gillies, S. Hagen and T. Schaub, The GeoJSON Format, RFC 7946, IETF, 2016. <https://www.rfc-editor.org/info/rfc7946>.
- [37] S. Das, S. Sundara and R. Cyganiak, R2RML: RDB to RDF Mapping Language, W3C Recommendation, W3C, 2012. <http://www.w3.org/TR/2012/REC-r2rml-20120927/>.
- [38] Schema.org. <http://schema.org/>.
- [39] M. Sporny, G. Kellogg and M. Lanthaler, JSON-LD 1.0. A JSON-based Serialization for Linked Data, W3C Recommendation, W3C, 2014. <http://www.w3.org/TR/2014/REC-json-ld-20140116/>.
- [40] W3C SPARQL Working Group, SPARQL 1.1 Overview, W3C Recommendation, W3C, 2013. <http://www.w3.org/TR/2013/REC-sparql11-overview-20130321/>.
- [41] F. Maali and J. Erickson, Data Catalog Vocabulary (DCAT), W3C Recommendation, W3C, 2014. <http://www.w3.org/TR/2014/REC-vocab-dcat-20140116/>.
- [42] ISA GeoDCAT-AP Working Group, GeoDCAT-AP: A geospatial extension for the DCAT application profile for data portals in Europe. Version 1.0.1, ISA Specification, European Commission, 2016. <https://joinup.ec.europa.eu/release/geodcat-ap/v101>.
- [43] R. Albertoni and A. Isaac, Data on the Web Best Practices: Data Quality Vocabulary, W3C Working Group Note, W3C, 2016. <https://www.w3.org/TR/2016/NOTE-vocab-dqv-20161215/>.
- [44] R. Battle and D. Kolas, Enabling the Geospatial Semantic Web with Parliament and GeoSPARQL, *Semantic Web Journal* **3**(4) (2012), 355–370. doi:10.3233/SW-2012-0065. <https://doi.org/10.3233/SW-2012-0065>.
- [45] S. Athanasiou, L. Bezati, G. Giannopoulos, K. Patroumpas and D. Skoutas, GeoKnow - Making the Web an Exploratory Place for Geo-spatial Knowledge. Market and Research Overview, Project Deliverable, GeoKnow FP7 Project, 2013. http://svn.aksw.org/projects/GeoKnow/Public/D2.1.1_Market_and_Research_Overview.pdf.
- [46] L. van den Brink, P. Janssen, W. Quak and J. Stoter, Linking spatial data: automated conversion of geo-information models and GML data to RDF, *International Journal of Spatial Data Infrastructures Research* **9** (2014), 59–85. doi:10.2902/1725-0463.2014.09.art3. <https://doi.org/10.2902/1725-0463.2014.09.art3>.
- [47] C.C. Aggarwal, N. Ashish and A.P. Sheth, The Internet of Things: A Survey from the Data-Centric Perspective, in: *Managing and Mining Sensor Data*, C.C. Aggarwal, ed., Springer, 2013, pp. 383–428. doi:10.1007/978-1-4614-6309-2_12. https://doi.org/10.1007/978-1-4614-6309-2_12.
- [48] U. Ramer, An iterative procedure for the polygonal approximation of plane curves, *Computer Graphics and Image Processing* **1**(3) (1972), 244–256. doi:10.1016/S0146-664X(72)80017-0. [https://doi.org/10.1016/S0146-664X\(72\)80017-0](https://doi.org/10.1016/S0146-664X(72)80017-0).
- [49] D.H. Douglas and T.K. Peucker, Algorithms for the Reduction of the Number of Points Required to Represent a Digitized Line or its Caricature, *The Canadian Cartographer* **10**(2) (1973), 112–122. doi:10.3138/FM57-6770-U75U-7727. <https://doi.org/10.3138/FM57-6770-U75U-7727>.
- [50] M. Visvalingam and J.D. Whyatt, Line generalisation by repeated elimination of points, *The Cartographic Journal* **30**(1) (1993), 46–51. doi:10.1179/000870493786962263. <https://doi.org/10.1179/000870493786962263>.
- [51] ISO/TC 211, ISO 19123:2005 Geographic information – Schema for coverage geometry and functions, ISO Standard, International Organization for Standardization (ISO), 2005. <https://www.iso.org/standard/40121.html>.
- [52] D. Brizhinev, S. Toyer and K. Taylor, Publishing and Using Earth Observation Data with the RDF Data Cube and the Discrete Global Grid System, W3C Working Group Note, W3C, 2017. <https://www.w3.org/TR/2017/NOTE-eo-qb-20170928/>.
- [53] M.K. Bergman, The Deep Web: Surfacing Hidden Value, White Paper, BrightPlanet, 2012. <https://brightplanet.com/2012/06/the-deep-web-surfacing-hidden-value/>.
- [54] B. He, M. Patel, Z. Zhang and K.C.-C. Chang, Accessing the Deep Web, *Commun. ACM* **50**(5) (2007), 94–101. doi:10.1145/1230819.1241670. <http://doi.acm.org/10.1145/1230819.1241670>.
- [55] J. Madhavan, D. Ko, Ł. Kot, V. Ganapathy, A. Rasmussen and A. Halevy, Google’s Deep Web Crawl, *Proc. VLDB Endow.* **1**(2) (2008), 1241–1252. doi:10.14778/1454159.1454163. <https://doi.org/10.14778/1454159.1454163>.
- [56] R.V. Guha, D. Brickley and S. Macbeth, Schema.org: Evolution of Structured Data on the Web, *Commun. ACM* **59**(2) (2016), 44–51. doi:10.1145/2844544. <https://doi.org/10.1145/2844544>.
- [57] C. McCathie Nevile and D. Brickley, HTML Microdata, W3C Working Draft, W3C, 2017. <https://www.w3.org/TR/2017/WD-microdata-20171010/>.
- [58] M. Sporny, HTML+RDFa 1.1 – Second Edition, W3C Recommendation, W3C, 2015. <http://www.w3.org/TR/2015/REC-html-rdfa-20150317/>.
- [59] J.R. Herring, OpenGIS® Implementation Standard for Geographic information - Simple feature access - Part 1: Common architecture, Version 1.2.1, OpenGIS® Implementation Standard, OGC, 2011. <http://www.opengeospatial.org/standards/sfa>.

- [60] G. Atemez and R. Troncy, Comparing vocabularies for representing geographical features and their geometry, in: *ISWC 2012, 11th International Semantic Web Conference, Terra Cognita Workshop, Volume 901*, Boston, US, 2012. <http://ceur-ws.org/Vol-901/paper1.pdf>.
- [61] ISA Action ARe3NA, Guidelines for the RDF encoding of spatial data, Technical specification, European Commission, 2017. <http://inspire-eu-rdf.github.io/inspire-rdf-guidelines>.
- [62] A. Perego, V. Ceti, A. Friis-Christensen and M. Lutz, GeoDCAT-AP: Representing geographic metadata by using the “DCAT application profile for data portals in Europe”, in: *Joint UNECE/UNGGIM Europe Workshop on Integrating Geospatial and Statistical Standards*, Stockholm, Sweden, 2017. <http://publications.jrc.ec.europa.eu/repository/handle/JRC107410>.
- [63] European Union, Establishing an Infrastructure for Spatial Information in the European Community (INSPIRE), Directive 2007/2/EC of the European Parliament and of the Council, 14 March 2007, *Official Journal of the European Union* **50**(L 108) (2007), 1–14. <http://data.europa.eu/eli/dir/2007/2/oj>.
- [64] A. Perego, A. Friis-Christensen and M. Lutz, GeoDCAT-AP: Use cases and open issues, in: *Smart Descriptions & Smarter Vocabularies (SDSVoc)*, Amsterdam, Netherlands, 2016. <http://publications.jrc.ec.europa.eu/repository/handle/JRC103804>.
- [65] I. Faniel, J. Pullmann and R. Atkinson, Dataset Exchange Use Cases and Requirements, W3C Working Group Note, W3C, 2017. <https://www.w3.org/TR/dcat-ucr/>.
- [66] R. Cyganiak and D. Reynolds, The RDF Data Cube vocabulary, W3C Recommendation, W3C, 2014. <http://www.w3.org/TR/2014/REC-vocab-data-cube-20140116/>.
- [67] L. Lefort, J. Bobruk, A. Haller, K. Taylor and A. Woolf, A Linked Sensor Data Cube for a 100 Year Homogenised Daily Temperature Dataset, in: *Proceedings of the 5th International Conference on Semantic Sensor Networks - Volume 904*, SSN’12, CEUR-WS.org, 2012, pp. 1–16. <http://ceur-ws.org/Vol-904/paper10.pdf>.
- [68] R. Atkinson, QB4ST: RDF Data Cube extensions for spatio-temporal components, W3C Working Group Note, W3C, 2017. <https://www.w3.org/TR/2017/NOTE-qb4st-20170928/>.
- [69] M. Jaber, M.A. Imran, R. Tafazolli and A. Tukmanov, 5G Backhaul Challenges and Emerging Research Directions: A Survey, *IEEE Access* **4** (2016), 1743–1766. doi:10.1109/ACCESS.2016.2556011. <https://doi.org/10.1109/ACCESS.2016.2556011>.
- [70] P. Barnaghi and A. Sheth, The Internet of Things: The Story So Far, *IEEE Internet of Things eNewsletter* (2014). <https://iot.ieee.org/newsletter/september-2014/the-internet-of-things-the-story-so-far.html>.
- [71] M. Simsek, A. Aijaz, M. Dohler, J. Sachs and G. Fettweis, 5G-Enabled Tactile Internet, *IEEE Journal on Selected Areas in Communications* **34**(3) (2016), 460–473. doi:10.1109/JSAC.2016.2525398. <https://doi.org/10.1109/JSAC.2016.2525398>.
- [72] J. Tomkins and D. Lowe, Timeseries Profile of Observations and Measurements. Version 1.0, OGC® Implementation Standard, OGC, 2016. <http://www.opengis.net/doc/IS/timeseries-profile-om/1.0>.
- [73] H. Hayashi, A. Asahara, K.-S. Kim, R. Shibasaki and N. Ishimaru, OGC Moving Features Access. Version 1.0, OGC® Implementation Standard, OGC, 2017. <http://www.opengis.net/doc/IS/movingfeatures-access/1.0>.
- [74] A. Aijaz, M. Simsek, M. Dohler and G. Fettweis, Shaping 5G for the Tactile Internet, in: *5G Mobile Communications*, W. Xiang, K. Zheng and X.S. Shen, eds, Springer, 2017, pp. 677–691. doi:10.1007/978-3-319-34208-5_25. https://doi.org/10.1007/978-3-319-34208-5_25.