Article

Bringing Federated Semantic Queries to the GIS-Based Scenario

Oswaldo Páez ¹ and Luis M. Vilches-Blázquez ^{2,*}

¹ Unidad Profesional Interdisciplinaria en Ingeniería y Tecnologías Avanzadas, Instituto Politécnico Nacional, La Laguna Ticoman, Mexico City 07340, Mexico; opaezo1500@alumno.ipn.mx

² Centro de Investigación en Computación, Instituto Politécnico Nacional, Av. Miguel Othón de Mendizabal s/n, UPALM—Zacatenco, Mexico City 07738, Mexico

* Correspondence: lmvilches@cic.ipn.mx

Abstract: Geospatial data is increasingly being made available on the Web as knowledge graphs using Linked Data principles. This entails adopting the best practices for publishing, retrieving, and using data, providing relevant initiatives that play a prominent role in the Web of Data. Despite the appropriate progress related to the amount of geospatial data available, knowledge graphs still face significant limitations in the GIScience community since their use, consumption, and exploitation are scarce, especially considering that just a few developments retrieve and consume geospatial knowledge graphs from within GIS. To overcome these limitations and address some critical challenges of GIScience, standards and specific best practices for publishing, retrieving, and using geospatial data on the Web have appeared. Nevertheless, there are few developments and experiences that support the possibility of expressing queries across diverse knowledge graphs to retrieve and process geospatial data from different and distributed sources. In this scenario, we present an approach to request, retrieve, and consume (geospatial) knowledge graphs available at diverse and distributed platforms, prototypically implemented on Apache Marmotta, supporting SPARQL 1.1 and GeoSPARQL standards. Moreover, our approach enables the consumption of geospatial knowledge graphs through a lightweight web application or QGIS. The potential of this work is shown with two examples that use GeoSPARQL-based knowledge graphs.

Keywords: GeoSPARQL; SPARQL; federated query; knowledge graph; geospatial data

Citation: Páez, O.; Vilches-Blázquez, L.M. Bringing Federated Semantic Queries to the GIS-Based Scenario. *ISPRS Int. J. Geo-Inf.* **2022**, *11*, 86. <https://doi.org/10.3390/ijgi11020086>

Academic Editors: Michail Vaitis and Wolfgang Kainz

Received: 24 November 2021

Accepted: 17 January 2022

Published: 25 January 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Geospatial data is increasingly being made available on the Web [1] in the form of knowledge graphs in the Semantic Web, often using Linked Data principles [2]. To achieve these knowledge graphs, (geospatial) resources need to be identified using HTTP URIs, indexed by search engines, and connected, or linked, to other resources [3]. Therefore, this entails adopting the best practices for publishing, retrieving, and using data on the Web [2,3]. These best practices are being embraced by a growing number of data providers, leading to building a global data space containing billions of assertions—the Web of Data [3].

The transformation and publication of geospatial data as knowledge graphs were pioneered by initiatives such as GeoNames (<http://www.geonames.org/ontology/documentation.html> (accessed on 23 November 2021)), OpenStreetMap [4], and Ordnance Survey [5]. After these initiatives, many geospatial datasets have been published in the Web of Data (<http://lod-cloud.net/> (accessed on 23 November 2021)). This has entailed geospatial data playing a pre-eminent role in the Web of Data cloud, operating as central nexuses that interconnect events, people, and objects [6] and offering an ever-growing semantic representation of the geospatial information wealth [7].

Despite the relevant progress related to the amount of data available, geospatial knowledge graphs still face significant limitations in the GIScience community since their use, consumption, and exploitation are really scarce in this community, especially considering

few existing works where these knowledge graphs are retrieved and integrated from within Geographic Information Systems (GIS).

To overcome these limitations and address some critical challenges of GIScience [8], standards (e.g., GeoSPARQL [9]) and domain-dependency best practices for publishing, retrieving, and using geospatial data on the Web have started appearing to propose this new perspective in the GIScience community. Additionally, some tools that combine the Semantic Web best practices and Web GIS have been collected in [10].

The (geospatial) knowledge graphs are stored and managed by triple stores, also known as RDF stores or knowledge bases. According to [1], they are capable of better addressing several types of issues at which relational databases struggle or are not intended to accomplish: queries with many joins across entities [11], queries with variable properties [11], and ontological inference on datasets. The World Wide Web Consortium (W3C) has collected a compendium of existing triple stores (https://www.w3.org/2001/sw/wiki/Category:Triple_Store (accessed on 23 November 2021)), where different implementations related to the geospatial domain can be found, and recent works have tested GeoSPARQL compliance in diverse triple store [7], highlighting, for instance, Apache Marmotta (<http://marmotta.apache.org/> (accessed on 23 November 2021)), Parliament (<https://github.com/SemWebCentral/parliament> (accessed on 23 November 2021)), OpenLink Virtuoso (<https://virtuoso.openlinksw.com/> (accessed on 23 November 2021)), or GeoSPARQL Fuseki (<https://jena.apache.org/documentation/geosparql/geosparql-fuseki> (accessed on 23 November 2021)). Nevertheless, the current developments do not support the possibility to express queries across diverse knowledge graphs. In this sense, there are not enough resources that support the combined implementation of GeoSPARQL and SPARQL 1.1 in order to allow performing federated geospatial queries on the Web of Data.

Considering this scenario and the recognized need for supporting federated queries using GeoSPARQL [12], we present an approach to request, retrieve, and consume (geospatial) knowledge graphs prototypically implemented on Apache Marmotta. This approach supports the combination of SPARQL 1.1 and GeoSPARQL standards, allowing performing federated geospatial queries on the Web of Data. Moreover, our approach enables consuming geospatial knowledge graphs displaying collected elements in a lightweight web application or adding them through an open-source GIS as QGIS. Additionally, our approach utilizes non-geospatial knowledge graphs to enrich the description of geospatial resources collected from knowledge graphs. The potential of our approach is shown with two examples that handle diverse GeoSPARQL-based knowledge graphs from the Web of Data.

In contrast to prior work, which uses the Linked Data paradigm from within a GIS [6], our approach focuses on a federated vision, is fully GeoSPARQL compliant, and is implemented on an open-source GIS such as QGIS. Furthermore, this work contributes to filling a gap and makes possible the exploitation of GeoSPARQL-based resources in a distributed environment through federated (queries) practices, helping the GIScience community to use the richness of the Web of Data.

The remainder of this paper is structured as follows. We motivate our work with a real-world scenario within the GIScience community in Section 2. Section 3 presents some concepts associated with background and some related works. A detailed description of methods and prototypical implementation of our federated GeoSPARQL queries approach is described in Section 4. Section 5 shows two examples of our approach's application using several (geospatial) knowledge graphs. Finally, discussion and conclusions are provided in Section 6.

2. Motivational Example

This work is motivated by the real-world scenario of the GIScience community where diverse users deal with layers overlapping as the geospatial data integrating procedure. In this scenario, the community often looks for data sources in various repositories, websites, or catalogs associated with Open Data or Spatial Data Infrastructures (SDIs) initiatives.

However, this state remains geospatial data as silos, making it challenging to achieve a whole and integrated vision of data available on the Web.

Several efforts have contributed to breaking down geospatial data silos by means of transforming “traditional” geospatial data into RDF [13–15], discovering links [15–18], exploiting topological relations [19,20], or combining the Semantic Web with SDIs [21–23]. Therefore, appropriate geospatial data sources can be used and queried in the Web of Data. Nevertheless, as we mentioned, these approaches focused on obtaining geospatial data out of data silos, whereas the use, consumption, and exploitation of knowledge graphs by the GIScience community remain scarce.

In this sense, we want our community to turn around knowledge graphs to use and consume existing geospatial data. Thus, we propose an approach to request, retrieve and consume geospatial knowledge graphs utilizing federated queries (SPARQL 1.1) and GeoSPARQL. Moreover, our approach allows for consuming geospatial knowledge graphs in a GIS since it provides potential beyond simply adding another data source [6] and contributes to overcoming semantic issues of GIS interoperability [24]. Thereby, it enables the application of modern spatial analysis’s immense toolboxes, as well as derives and exploits knowledge from heterogeneous data sources. In this way, our approach allows users to manipulate geospatial knowledge graphs in two ways: (i) exploiting data directly through federated GeoSPARQL queries and displaying collected data in an application web, and (ii) consuming geospatial knowledge graphs through a GIS to exploit this information in a familiar environment for the community.

At first glance, the scenario mentioned above appears as nothing more than a software engineering task, namely, an improved import functionality for GIS that would add geospatial knowledge graphs to the already extensive set of data formats that can be consumed. However, knowledge graphs are not a data format. As we mentioned before, it is a paradigm based on Linked Data principles with strategies needed for Web-scale, distributed data infrastructure, and does not adapt well with how we exchange data in GIS [6]. Therefore, it entails new ways of thinking about how GIScience and its community should interact with (geospatial) knowledge graphs in a distributed environment and through federated (queries) ways, allowing us to exploit the richness of a global data space containing billions of (geospatial) assertions—the Web of Data [3].

3. Background and Related Work

3.1. Knowledge Graphs

The *Knowledge Graph* concept, also recognized as Semantic Networks, was used as early as the 1970s to represent human knowledge in a machine-readable form [25,26]. Later, Google rediscovered and spread a new vision of this concept when it introduced its notion about a new web search strategy in 2012 (<https://googleblog.blogspot.com/2012/05/introducing-knowledge-graph-things-not.html> (accessed on 23 November 2021)). It involved changing from pure text processing to a more symbolic representation of knowledge, expressed in the following way:

“The Knowledge Graph enables you to search for things, people or places that Google knows about landmarks, celebrities, cities, sports teams, buildings, geographical features, movies, celestial objects, works of art and more and instantly get information that is relevant to your query. This is a critical first step towards building the next generation of search, which taps into the collective intelligence of the web and understands the world a bit more like people do.”

The Semantic Web context has also promoted a graph-based representation associated with the Resource Description Framework (<https://www.w3.org/RDF/> (accessed on 23 November 2021)) (RDF), which is the standard knowledge representation language for this Web. RDF provides a reliable infrastructure to publish, share, and query structured data on the Semantic Web [27]. In addition, RDF gives a framework for describing information on the web, of which the syntax (a data model) is two-fold: (1) RDF graphs are sets of *Subject-Predicate-Object* triples where the components may be URIs, datatyped literals, or

blank nodes. The graphs present resource descriptions, for instance: metadata about social networks, digital artifacts, personal information, etc., and also supply a mechanism of integration over heterogeneous information sources; (2) RDF datasets are employed to make RDF graph collections and generate a default graph and zero or more named graphs.

According to [28], from a broad panorama, any graph-based representation of some knowledge could be considered a knowledge graph since no universal definition exists regarding what a knowledge graph is and what it is not [29]. However, we select the minimum set of components of knowledge graphs suggested by [28]:

- Mainly describes real-world entities and their interrelations, organized in a labeled directed graph.
- Defines possible classes and relations of entities in a schema.
- Allows for potentially interrelating arbitrary entities with each other.
- Covers various topical domains.
- Allows us to infer data through relationships among nodes.

Numerous examples of knowledge graphs in the Semantic Web are often available using the Linked Data principles [2]. These principles are characterized by: (1) the use of URIs as names for things; (2) the use of HTTP URIs so that people can look up those names; (3) when someone looks up a URI, giving valuable information through standards Resource Description Framework (RDF) and SPARQL Protocol and RDF Query Language (<https://www.w3.org/TR/sparql11-query/> (accessed on 23 November 2021)) (SPARQL); and (4) adding links to other URIs so that they can find more things. Some of these knowledge graphs are DBpedia (<https://dbpedia.org/> (accessed on 23 November 2021)), OpenCyc (<https://github.com/asanchez75/opencyc> (accessed on 23 November 2021)), Wikidata (<https://www.wikidata.org> (accessed on 23 November 2021)), Cyc (<https://www.cyc.com/> (accessed on 23 November 2021)), etc. A comprehensive description of these knowledge graphs and other ones is described in [28].

3.2. SPARQL and GeoSPARQL

SPARQL is a semantic query language able to retrieve and manipulate data stored in RDF. This language can express queries across diverse data sources, whether data are stored natively as RDF or viewed as RDF via middleware. In addition, SPARQL contains capabilities for querying required and optional graph patterns along with their conjunctions and disjunctions. SPARQL also supports aggregation, subqueries, negation, creating values by expressions, extensible value testing, and constraining queries by source RDF graph. The results of SPARQL queries can be result sets or RDF graphs.

The development of SPARQL 1.1 enables expression queries across diverse knowledge graphs, whether data are stored natively as RDF or viewed as RDF via middleware. The W3C specification establishes the syntax and semantics of SPARQL 1.1 Federated query extension for executing queries distributed over several SPARQL Endpoints. Additional details about this standard can be found in the W3C Specifications <https://www.w3.org/TR/sparql11-federated-query/> (accessed on 23 November 2021).

Open Geospatial Consortium (OGC) has defined an approach for describing and querying geospatial data on the Semantic Web. This effort, called GeoSPARQL [9], defines a vocabulary for representing geospatial data in RDF and an extension to the SPARQL query language for processing this type of data. In addition, GeoSPARQL is designed to accommodate systems based on qualitative spatial reasoning and systems based on quantitative spatial computations. This proposal enables defining distinct kinds of geometry (e.g., polygons, lines, points, multipoints, etc.), adding multiple coordinate reference systems, and bringing to the Web of Data cloud the chance to reveal spatial relations to query geographic datasets (e.g., touches, intersects, overlaps, etc.). The class *Geometry* represents geometries and coordinates can be encoded in an RDF literal of type Well-Known Text (WKT), which is associated with [30], utilizing a single RDF property, namely *asWKT*.

GeoSPARQL also allows for encoding geometries using the Geography Markup Language (GML). Therefore, the data type (*GMLLiteral*), property (*asGML*), and the URL

for the geometry type have to be adjusted accordingly. Further specifications regarding GeoSPARQL are described in [9].

3.3. Related Work

Several works have put knowledge graphs from the Semantic Web perspective closer to the GIScience community. In this context, some approaches concentrated on the transformation process from geospatial data to RDF. Schade and Cox [31] presented a proposal to convert GML to RDF and extend data types given by WFS to other types, such as GML, RDF, and HTML. Stadler et al. [32] built an ETL approach to transforming OpenStreetMap data to Linked Data, creating the LinkedGeoData initiative. Usery and Varanka [13] defined a conversion from vector and raster data to RDF and also encoded vector data in GML format utilizing pre-computing topological relations and converted them into RDF triples. van den Brink et al. [14] described a semi-automatic transformation from geographic information models and GML data to RDF data. In addition, different tools have emerged to perform the transformation of “traditional” geospatial data (shapefiles or geospatial databases) into RDF a little easier, such as Geometry2RDF [33], Sparqlify, TripleGeo [34], SHP2GeoSPARQL [23], or GeoTriples [35].

Other approaches aimed at a more profound integration by enabling spatial query processing within SPARQL [10,36,37] by implementing a Linked Data connector framework as a set of toolboxes for Esri’s ArcGIS to allow the retrieval, integration, and analysis of Linked Data from within GIS [6], or by developing new standards or extensions for spatial and temporal reasoning and querying over Linked Data [38].

On the other hand, diverse works focused on handling geospatial data using GeoSPARQL. Battle and Kolas [1] developed one of the first GeoSPARQL implementations, enabling buffer queries for nearby places such as parks. Bereta et al. [39] described Ontop-spatial, a geospatial extension of the OBDA system Ontop, that leverages the technologies of geospatial databases and facilitates GeoSPARQL-to-SQL translation. The authors depicted the system’s functionalities discovering agricultural fields that intersect with protected areas and integrating information regarding floods. Nys et al. [40] employed GeoSPARQL and OWL-Time to extend CIDOC CRM: CRMgeo, providing a near-full data management for complex spatio-temporal reasoning and querying through SPARQL queries. Vilches-Blázquez and Saavedra [23] performed a link discovery process based on a spatial analysis process using GeoSPARQL. This work allowed for setting spatial matchings between different geospatial data sources specifying topological relations (e.g., contain, intersect, touch, etc.) based on GeoSPARQL.

With respect to federated queries, some works have employed federated queries on (geospatial) knowledge graphs [41–46]. However, the application of queries with geospatial functions is limited, and GeoSPARQL is not entirely compliant. Additionally, considering the increasing approaches based on GeoSPARQL, some works have provided ways to measure the support in GeoSPARQL-enabled RDF triple stores [7,47–49]. Even a benchmark utilizing GeoSPARQL constructs was defined, facing all phases of federated query processing [50].

To sum up, several approaches have focused mainly on adding the geospatial component to the Web of Data, developing relevant use cases and tools for transforming “traditional” geospatial data in knowledge graphs (even based on GeoSPARQL), and providing some spatial capabilities. In addition, few works have dealt with the retrieval and integration of knowledge graphs from within GIS. However, to the best of our knowledge, no previous results exist where geospatial knowledge graphs are requested and retrieval through a federated way combining SPARQL 1.1 and GeoSPARQL standards. Moreover, although a related work exists where Linked Data is retrieved and integrated within ArcGIS [6], our approach focuses on a federated vision, utilizes fully GeoSPARQL compliant, and is implemented on an open-source GIS such as QGIS.

4. Methods and Prototypical Implementation

This section presents a brief description of the triple store on which we developed the federated geospatial approach's reference implementation (Apache Marmotta) and describes the main characteristics related to the architecture and design of our approach, detailing each of its components.

4.1. Apache Marmotta

Apache Marmotta is an open-source triple store that, in 2012, joined the Apache Software Foundation incubator. The Apache Software Foundation governs it and, in November 2013, it was graduated as a Top-Level Project under the Apache Software License version 2.0. Apache Marmotta started as a Linked Media Framework project whose main goal was to be an easy-to-set-up server application that joined some technologies such as Apache Stanbol and Apache Solr. Apache Marmotta offers a platform to read, write, and update RDF data through ontologies, graph databases, and the SPARQL language.

It can be deployed on Windows through a .jar file and executed as such file. The installation of Apache Marmotta using this method makes its deployment easier since users do not need to configure any setting, including the Apache Tomcat webserver. Another way to install this triple store is to compile source code with Maven and execute it or even run it with a public Docker image.

Concerning data persistence, RDF data can be stored in three different databases: H2 (by default), MySQL, whose connector cannot be distributed, and PostgreSQL needed for geoprocessing tasks. Taking into account the scope of this work, PostgreSQL and its geospatial extension, PostGIS, were set up and used in Apache Marmotta.

With respect to the considered standards of this work, Apache Marmotta supports an appropriate number of features related to SPARQL 1.1 and offers the possibility of processing geometries (e.g., points, lines, polygons, and multipolygons) with 35 functions that support GeoSPARQL. However, both standards are not integrated into Apache Marmotta, coexisting as isolated components. Therefore, this triple store cannot process geospatial resources in a federated way.

4.2. Architecture and Design

This approach focuses on the challenge to perform federated queries combining SPARQL 1.1 and GeoSPARQL in the context of the geospatial knowledge graphs. Doing so, we present different steps and implementation for the triple store Apache Marmotta. Figure 1 depicts our workflow to carry out federated geospatial queries within this *new data ecosystem* for the GIScience community. Thereby, our approach starts gathering geospatial features from different (1..n) SPARQL Endpoints, for instance, Find all features (of 1 or n types) from a selected location (s). It allows for collecting geospatial features from knowledge graphs that support GeoSPARQL. Next, using gathered geospatial features as the starting point, our approach utilizes spatial operators associated with GeoSPARQL to discover new features from diverse (1..n) knowledge graphs through their related SPARQL Endpoints. The following step helps filter geospatial features collected from distinct knowledge graphs to handle easily obtained data conforming to user preferences. The fourth step allows for enriching geospatial features with additional information from other repositories (e.g., DBpedia or Wikidata). Finally, our approach enables us to handle and exploit elements collected from diverse knowledge graphs through an application web or GIS application, providing a new scenario to deal with federated geospatial information in the Web of Data. In order to achieve these steps of our approach, some issues were overcome, which are described in the following sections.

Figure 1. The workflow of the geospatial federated queries connector.

The obtained results through various steps of our federated GeoSPARQL approach can be stored into a local Apache Marmotta deployment, allowing a materialization of the elements collected from knowledge graphs of the Web of Data. It enables us to query and exploit (geospatial) federated resources in a local environment, simulating the traditional operation of GIS applications. Likewise, our approach allows for exploiting the mentioned resources directly using federated GeoSPARQL queries to knowledge graphs.

4.2.1. Gathering Geospatial Knowledge Graphs

As we mentioned before, Apache Marmotta cannot handle federated geospatial resources. Considering this limitation, our approach allows gathering diverse geospatial knowledge graphs through federated connection to REST services associated with SPARQL Endpoints of the Web of Data. In this way, the system performed an extension of the Apache Marmotta software to support the connection of SPARQL 1.1 and GeoSPARQL. This extension handles unexpected exceptions in GeoSPARQL classes and manipulates data available within the Apache Marmotta in order to parse data to PostgreSQL/PostGIS dialect for processing collected data. Thus, a key component of this step focuses on preprocessing GeoSPARQL data, sending the related query to PostgreSQL, and returning data to Apache Marmotta based on SPARQL 1.1.

This component of our approach is based on Listing 1, which inputs a URL for each SPARQL Endpoint and includes a SPARQL query to collect geospatial features (*geoFeature*) from distinct knowledge graphs of the Web of Data. These components allow for retrieving information about *geoFeature*, by default, URI, label, and geometry. The next step of this query plan is to retrieve elements of each knowledge graph from each SPARQL Endpoint.

Listing 1. Query plan to gather *geoFeature* from knowledge graphs.

```

Require: SPARQL Endpoint's URLs, geoFeature, geoFeatureType, and geometry
Ensure: Gathering geoFeature from knowledge graphs  $w(\text{URI, label, geometry})$ 
 $w \leftarrow \{\}$ ;
 $w.\text{URI} \leftarrow \text{URI}$ 
{Extract geoFeature from specific knowledge graphs (KGs)}
SPARQLEndpoint  $\leftarrow \text{URI } geoFeature$ ,
For all  $geoFeature \in \text{SPARQLEndpoint}$  do
Set  $geoFeature = \text{query SPARQLEndpoint}_1 ("geoFeatureType")$ 
UNION
query  $\text{SPARQLEndpoint}_2 ("geoFeatureType")$ 
UNION
query  $\text{SPARQLEndpoint}_n ("geoFeatureType")$ 
 $geoFeatureList \leftarrow geoFeature(\text{URI, label, geometry})$ 
end for

```

4.2.2. Applying GeoSPARQL Functions and Filtering Features on a Federated Scenario

This component of our approach is especially relevant since it supports querying and exploiting geospatial knowledge graphs based on GeoSPARQL. It is possible since our development supports the 35 GeoSPARQL functions to exploit these knowledge graphs in a federated way within the Web of Data scenario. Additionally, our approach can combine geospatial operators and filters to obtain specific resources from the Web of Data in a federated way. Thus, the system can utilize the previous query results to perform a spatial query (based on GeoSPARQL) using a set of properties that qualitatively represent the spatial relations defined in RCC-8 (Region Connection Calculus). This process can enrich original resources with additional data collected from distributed (Geo)SPARQL Endpoints, and considering the data richness presented in the Web of Data, our approach allows filtering resources obtained from spatial queries.

This component of our approach is based on Listing 2, which takes information about *geoFeature(s)* collected from the previous step. Considering these *geoFeature(s)*, our approach retrieves additional *geoFeature(s)* from the same or extra knowledge graphs. To gather this/these new *geoFeature(s)*, our approach applies spatial operations based on RCC8 since our approach is fully GeoSPARQL-complaint. In addition, this step allows for filtering *geoFeature(s)* using one or multiple types of *geoFeature(s)*, increasing the potential of the geospatial query. The next stage of this query plan is to collect elements of each knowledge graph from each SPARQL Endpoint.

Listing 2. Query plan to apply GeoSPARQL and filters by features.

```

Require: SPARQL Endpoint's URLs, geoFeature, geoFeatureType, geometry, and RCC8 operator(s).
Ensure: Gathering geoFeature from knowledge graphs  $w(\text{URI, label, geometry})$  after applying spatial operator(s)
For  $geoFeatureList$  do
Set  $geoFeature' = \text{query SPARQLEndpoint}_1 ("geoFeatureType_1")$ 
UNION
query  $\text{SPARQLEndpoint}_2 ("geoFeatureType_2")$ 
UNION
query  $\text{SPARQLEndpoint}_n ("geoFeatureType_n")$ 
 $geoFeatureList' \leftarrow geoFeature'(\text{URI, label, geometry})$ 
 $geoFeatureListRCC8 \leftarrow \text{RCC8\_operator}(geoFeatureList'), \text{geometry,}$ 
 $\text{dataType(boolean or float)}$ .
end for

```

4.2.3. Enriching with Additional Information

Any data source often includes specific references to other data types, and geographic data sources are no anomaly. These references extend their significance when we deal with

expert domain data sources. Nevertheless, non-expert users often find it challenging to know and apply these hidden data [51]. Therefore, the spatial enrichment process proposes recovering such information and making it explicit [52].

Our approach enables us to collect additional information from geospatial features extracted from the combination of geospatial operators and filters in the previous step to enrich the description of specific resources with supplementary information from the Web of Data. We consider that this workflow step can enhance the features since geospatial data are sometimes published with limited attributes focused on geometrical characteristics. Therefore, it can be helpful to utilize the Web of Data to enrich (geospatial) resources with more descriptive information. Nevertheless, one aspect to consider of our approach is that we cannot always guarantee an enrichment with additional information since, in some cases, our proposal cannot establish any match between (geospatial) resources and the Web of Data's resources.

In order to perform this enrichment process, the approach relies on resource names (*rdfs:label*) and types (*rdf:type*) to look for these resources in the Web of Data, concretely, in DBpedia, although the approach provides a flexible setting to define any SPARQL Endpoint that users consider. DBpedia is a critical resource of this process because it is regarded as a cross-domain knowledge base and remains a nucleus of the Web of Data [53]. Its ontology collects 685 classes that form a subsumption hierarchy and are described by 2795 properties. Furthermore, this ontology comprises approximately 4,233,000 instances, from which 735,000 instances are associated with the Place class. In this way, this knowledge graph may be a key element to enrich any geospatial resource.

Listing 3. Query plan to collect additional information.

```

Require: DBpedia SPARQL Endpoint's URLs, geoFeature, geoFeatureType, label.
Ensure: Gathering additional information about geoFeature(s) from DBpedia z(URI, subject,
abstract, type, and image)
z ← {};
z.URI ← URI
{Extract additional data related to geoFeature from DBpedia }
geoFeatureList' ← Label geoFeature',
For all geoFeature' ∈ SPARQLEndpoint do
Set geoFeature' = query DBpediaSPARQLEndpoint1 ("label")
OPTIONAL
query DBpediaSPARQLEndpoint1 ("abstract")
OPTIONAL
query SPARQLEndpointn ("subject")
OPTIONAL
query SPARQLEndpointn ("type")
OPTIONAL
query SPARQLEndpointn ("image")
geoFeatureList'' ← geoFeature' (URI, subject, abstract, type, and image)
end for

```

Listing 3 presents a SPARQL query plan to collect additional information from geospatial features. Although we think this step can be helpful to gather information regarding features filtered spatially and by type (s) (from the previous step), this query plan can be performed about any features collected in the distinct phases of our approach. Thus, the query plan uses type and label elements of the *geoFeature(s)* collected previously to search these features in DBpedia. When a match is carried out, our approach recovers additional information associated with abstract, subject, label (in different languages, when are available), and images. Due to these elements can be omitted in some cases (*geoFeature*), they are considered *OPTIONAL* elements in the query plan. The next stage of this query plan is to collect elements from DBpedia and offer them to the user enriching the original data.

4.2.4. Exploiting Results

To explore the results obtained from previous steps, we developed two options for dealing with results obtained from federated GeoSPARQL queries. On the one hand, a web application that works as a SPARQL client and shows geospatial data on a map. The web application was developed with Node.js and Leaflet, and it provides end-users a SPARQL Endpoint connected to Apache Marmotta for writing federated GeoSPARQL queries. When a query is executed (we assume that it is formed correctly), the application allows for presenting results of the query as JSON or a table. In addition, when the results contain geospatial information, a third option is offered, enabling displaying information based on GeoSPARQL on a map.

On the other hand, we connected our approach to the QGIS tool since we attempted to close (geospatial) knowledge graphs to the GIScience community. Thus, we reused the SPARQLing Unicorn QGIS Plugin (<https://plugins.qgis.org/plugins/sparqlunicorn/> (accessed on 23 November 2021)) to connect to our approach and QGIS. It allows adding GeoJSON layers to QGIS with data collected from GeoSPARQL federated queries on the Web of Data. Herein, the plugin requires connection parameters to the SPARQL Endpoint of the selected knowledge graphs to collect existing features and geometries (see Figure 2). A detailed explanation of connecting our approach and QGIS is available on <https://github.com/Osw1997/Guide-connection-for-Apache-marmotta-and-QGIS> (accessed on 23 November 2021).

Figure 2. The SPARQLing Unicorn QGIS Plugin to obtain features from (geospatial) knowledge graphs.

5. Examples

In this section, we present two examples of using our federated GeoSPARQL queries approach using diverse resources of the Web of Data. Next, we provide related details to each one of them.

5.1. Bicycle Parking and Historical POIs

In this running example, we dealt with four distinct knowledge graphs such as *datos.zaragoza.es* (a knowledge graph with open data from the Zaragoza city), Linked-GeoData (an effort that uses the information collected by the OpenStreetMap project and

makes it available as an RDF knowledge base according to the Linked Data principles—<http://linkedgeodata.org/> (accessed on 23 November 2021)), [datos.ign.es](https://datos.ign.es/btn100.html) (accessed on 23 November 2021) (A knowledge graph that provides information related to the National Topographic Database (*Base Topográfica Nacional* 1:100.000—BTN100) (<https://datos.ign.es/btn100.html> (accessed on 23 November 2021)) from the Spanish National Mapping Agency (*Instituto Geográfico Nacional de España* (IGN-E)—<http://www.ign.es/> (accessed on 23 November 2021)), and a local repository with resources deployed from LinkedGeoData.

This example looks for all information on bicycle parking from a specific location, concretely in Zaragoza (Spain). For that, we constructed a SPARQL query to collect bicycle parking of this city in several knowledge graphs using the federated component of our system. Listing 4 depicts a SPARQL query submitted to two different repositories, concretely to <https://www.zaragoza.es/sede/portal/datos-abiertos/servicio/sparql> (accessed on 23 November 2021) and LinkedGeoData, to obtain bicycle parking of our study case.

Listing 4. A SPARQL query to datos.zaragoza.es and LinkedGeoData.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX geof: <http://www.opengis.net/def/function/geosparql/>
PREFIX units: <http://www.opengis.net/def/uom/OGC/1.0/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX vocab: <http://vocab.linkeddata.es/datosabiertos/def/urbanismo-infraestructuras/>

select * where {
  # Get geometries associated with bicycle parking in Spain (Zaragoza's SPARQL endpoint).
  service <http://datos.zaragoza.es/sparql> {
    select ?sb ?park where {
      ?sb ?ob vocab:equipamiento#AparcamientoBicicleta.
      ?sb ?p ?o ;
      <http://www.w3.org/2003/01/geo/wgs84_pos#geometry> ?o .
      ?o <http://www.opengis.net/ont/geosparql#asWKT> ?og .
      # Filter data that has the geodetic system WGS84
      filter(regex(?o, "WGS84"))
      bind(str(?og) as ?park)
    } limit 50
  }
}
```

The following step of our approach allows applying geoprocessing using GeoSPARQL spatial operators about geospatial data collected previously. Therefore, we utilized bicycle parking from Zaragoza gathered from <https://www.zaragoza.es/sede/portal/datos-abiertos/servicio/sparql> (accessed on 23 November 2021) and LinkedGeoData and created a buffer around each bicycle parking to recover near Points of Interest (POIs) from various knowledge graphs of the Web of Data. Moreover, if we are interested in specific POIs, our approach allows for combining geospatial operators and filters to obtain particular POIs. Listing 5 presents a GeoSPARQL federated query to obtain historical POIs near bicycle parking collected by the query of Listing 4. In this case, (Listing 5), the GeoSPARQL federated query is submitted to four distinct SPARQL Endpoints considered, such as <https://www.zaragoza.es/sede/portal/datos-abiertos/servicio/sparql> (accessed on 23 November 2021), LinkedGeoData, datos.ign.es (accessed on 23 November 2021), and a local LinkedGeoData repository. In addition, due to the combination of geospatial operators and filters by POI type applied into Listing 5, we obtained some of the results shown in Table 1, where diverse bicycle parking has associated near historical POIs and distance among them in meters (1000 m).

Table 1. A results excerpt of the spatial operators and filters combination.

BicyclePark	Geometry	KGs	POI	Geometry	Distance_m
lgd:node1256882341	POINT (−0.8823273 41.6531583)	LGD	La Audiencia o Palacio de los Condes de Luna	POINT (−0.88220066 41.652898609535)	30.71180431220672
lgd:node1256882341	POINT (−0.8823273 41.6531583)	LGD	Palacio de los Condes de Sástago	POINT (−0.88199995 41.652768769535)	51.14054869447922
lgd:node1256882341	POINT (−0.8823273 41.6531583)	LGD	Palacio de Fuenclara	POINT (−0.88223207 41.653999079536)	93.71972830849011
lgd:node1256882341	POINT (−0.8823273 41.6531583)	LGD	Monumento a los mártires de la religión y de la patria	POINT (−0.8810072000 41.6520520000)	164.8951479494759
zgz:198	POINT (−0.875644 41.659145)	ZGZ	Puente de Piedra y Pretíl de San Lázaro	POINT (−0.87536409 41.6571106695)	227.1480092502848
lgd:node1256882341	POINT (−0.8823273 41.6531583)	LGD	Palacio de los Condes de Argillo	POINT (−0.88304861 41.6552085795)	235.5130207000509
zgz:100	POINT (−0.873681 41.660137)	ZGZ	Puente de Piedra y Pretíl de San Lázaro	POINT (−0.87536409 41.6571106695)	364.1904928910845
...	...	...	...	...	...

@ prefix lgd: <http://linkedgeodata.org/triplify/> (accessed on 23 November 2021). @ prefix zgz: <http://www.zaragoza.es/api/recurso/urbanismo-infraestructuras/equipamiento/aparcamiento-bicicleta> (accessed on 23 November 2021).

Listing 5. A GeoSPARQL federated query to obtain historical POIs near bicycle parking.

```

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX geof: <http://www.opengis.net/def/function/geosparql/>
PREFIX units: <http://www.opengis.net/def/uom/OGC/1.0/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX vocab: <http://vocab.linkeddata.es/datosabiertos/def/urbanismo-infraestructuras/>

select * where {
  # Get geometries associated to bicycle parking in Spain (Zaragoza's SPARQL endpoint).
  service <http://datos.zaragoza.es/sparql> {
    select ?sb ?park where {
      ?sb ?ob vocab:equipamiento#AparcamientoBicicleta .
      ?sb ?p ?o ;
      <http://www.w3.org/2003/01/geo/wgs84_pos#geometry> ?o .
      ?o <http://www.opengis.net/ont/geosparql#asWKT> ?og .
      # Filter data that has the geodetic system WGS84
      filter(regex(?o, "WGS84"))
      bind(str(?og) as ?park)
    } limit 50
  }

  # Recover Points Of Interest (POI) from IGN.es
  service <https://datos.ign.es/sparql/> {
    select ?titlePOI ?geometriaPOI where {
      ?s a <https://datos.ign.es/def/btn100#LugarDeInteres> .
      ?s <http://purl.org/dc/terms/title> ?titlePOI .
      ?s <http://www.opengis.net/ont/geosparql#hasGeometry> ?geoIns .
      ?geoIns <http://www.opengis.net/ont/geosparql#asWKT> ?geom .
      bind(str(?geom) as ?geometriaPOI)
      bind("POI" as ?tipo)
    } limit 25 offset 100
  }

  # Recover historic places from our local server that contains a dataset from LinkedGeoData.
  service <http://192.168.100.10:7200/repositories/IDRep> {
    select ?titleHistoric ?geometriaHistoric where {
      graph <http://www.linkedGeoData.org/historic> {
        ?nodes <http://www.w3.org/2000/01/rdf-schema#label> ?labelEsp_ .
        ?nodes <http://geovocab.org/geometry_geometry> ?geoIns .
        ?geoIns <http://www.opengis.net/ont/geosparql#asWKT> ?geoEsp_ .
        bind(str(?geoEsp_) as ?geometriaHistoric) .
        bind(str(?labelEsp_) as ?titleHistoric) .
        bind("TurismoEspania" as ?tipo)
      }
    }
  }

  # Then, get distance for each combination of points
  # from 3 previous graphs (Bicycles, POIs and historic places).
  bind(geof:distance(?park, ?geometriaPOI, units:meter) as ?D_m_POI)
  bind(geof:distance(?park, ?geometriaHistoric, units:meter) as ?D_m_Historic)
  # Finally, exclude points where distance is greater than one kilometer between points.
  filter((?D_m_POI < 1000) || (?D_m_Historic < 1000))
}

```

Continuing with the workflow of our approach, we collected additional information about historical POIs (obtained during the last step) near bicycle parking. Thus, for instance, we enriched the original data of the Palacio de Fuenclara (https://es.dbpedia.org/page/Palacio_de_Fuenclara (accessed on 23 November 2021)) with the description (dbo:abstract), image (prop-es:imagen) and related subjects (dct:subject) provided by esDBpedia (see Table 2). For this case, we used the Spanish DBpedia version (<https://es.dbpedia.org/>

(accessed on 23 November 2021)) because we dealt with information related to Spain, and, hence, this specific DBpedia version has more detailed resources that may contribute to enriching geospatial resources collected from the previous steps of our approach.

Table 2. Results from DBpedia.

POI	Geometry	Subject	Image	Abstract
Palacio de Fuenclara	POINT (-0.88223207 41.653999079536)	dbpedia:Category% C3%ADa:Palacios _de_Zaragoza	dbpedia:Archivo:Palacio _de_Fuenclara.jpg	<i>El palacio de Fuenclara en Zaragoza (España) fue construido en la segunda mitad del siglo XVI por encargo de don Antonio Agustín, padre del arzobispo de Tarragona y eminente canonista, transformado en el siglo XVII por sus nuevos inquilinos, los Condes de Fuenclara . . .</i>

@ prefix dbpedia: <<http://es.dbpedia.org/resource/>> (accessed on 23 November 2021).

As mentioned above, our approach supports different ways to display the gathered results. In this case, we visualized the collected elements from various geospatial knowledge graphs on the developed web application (see Figure 3).

Figure 3. Results displayed on the web application.

5.2. Zoos and Museums

This example considers three different knowledge graphs: LinkedGeoData, datos.ign.es (accessed on 23 November 2021), and FOODIE (a cloud platform for gathering, storing, sharing, and analyzing spatially and non-spatially referenced data related to agriculture, forestry, and environment—<https://www.foodie-cloud.org/> (accessed on 23 November 2021)).

Within this example, we look for zoos and museums contained into a specific administrative boundary in Spain, concretely, provinces. To gather these resources from the Web of Data, we constructed a GeoSPARQL query to obtain all the geometries of different provinces and complete the mentioned query to recover all zoos available in Spain (from Foodie's SPARQL Endpoint) and museums (from LinkedGeoData's SPARQL Endpoint). After obtaining the geometries, we included the GeoSPARQL operation “contains”, whose input parameters are the variables associated with collected geometries. Listing 6 shows the query executed in Apache Marmotta to gather the resources of this study case from the Web of Data.

In this case, we do not provide the results in a table because the geometries of each province offer literals with a detailed collection of coordinates. After all, these regions are composed of complex geometry types. Regarding visualization of the results, we set the previous query (see Listing 6) in QGIS using the extension of SPARQLing Unicorn QGIS Plugin to support our approach developed on Apache Marmotta. Figure 4 depicts results collected, where provinces are presented as polygons (blue shapes), zoos are symbolized as red dots, and museums as green dots.

Figure 4. Results displayed on QGIS.

Listing 6. A GeoSPARQL federated query to obtain zoos and museums contained in Spanish provinces.

```

PREFIX geof: <http://www.opengis.net/def/function/geosparql/>
PREFIX units: <http://www.opengis.net/def/uom/OGC/1.0/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX vocab2: <<http://vocab.linkeddata.es/datosabiertos/def/sector-publico/>

select ?titleProv ?titlePlace ?placeFrom ?geoProv ?geoPlace where {
  # Select Provinces
  service <https://datos.ign.es/sparql> {
    select ?s ?titleProv ?geoProv where {
      ?s ?p vocab2:territorio#Provincia .
      ?s <http://www.opengis.net/ont/geosparql#hasGeometry> ?o ;
      # Province's name
      <http://purl.org/dc/terms/title> ?titleProv .
      ?o <http://www.opengis.net/ont/geosparql#asWKT> ?geo_ .
      bind(str(?geo_) as ?geoProv)
    }
  }

  # Make an UNION between Zoos and museums in order to create a bigger graph.
  {
    # Zoos
    service <https://www.foodie-cloud.org/sparql> {
      select ?titlePlace ?geoPlace ?placeFrom where {
        ?sF a <http://gis.zcu.cz/SPOI/Ontology#zoo> . ?sF <http://purl.org/dc/elements/1.1/title>
        ?titlePlace .
        ?sF <http://www.opengis.net/ont/geosparql#asWKT> ?geo_ .
        ?sF <http://www.opengis.net/ont/geosparql#sfWithin> ?w .
        bind(str(?geo_) as ?geoPlace)
        bind("Foodie" as ?placeFrom)
        filter(regex(?w, "Spain")) .
      } limit 25
    }

    UNION {
      # Museums
      service <http://linkedgeodata.org/sparql> {
        select distinct ?titlePlace ?geoPlace ?placeFrom where {
          ?sb ?p <http://linkedgeodata.org/ontology/Museum> .
          ?s <http://www.w3.org/2000/01/rdf-schema#label> ?titlePlace .
          ?s <http://geovocab.org/geometry#geometry> ?o .
          ?o <http://www.opengis.net/ont/geosparql#asWKT> ?o2 .
          bind(str(?o2) as ?geoPlace)
          bind("LinkedGeoData" as ?placeFrom)
          bind(xsd:float(strbefore(strafter(str(?geoPlace), " "), "")) as ?c2)
          bind(xsd:float(strbefore(strafter(str(?geoPlace), "("), "")) as ?c1)
          filter(?c1 > -1 && ?c1 < 0 && ?c2 > 40 && ?c2 < 42)
        } limit 25
      }
    }

    # Filter geometries where provinces contains at least one zoo/museum.
    bind(geof:sfContains(?geoProv, ?geoPlace) as ?contains)
    filter(str(?contains) = 't')
  }
}

```

6. Discussion and Conclusions

As we mentioned in this work, several approaches have focused on adding the geospatial component to the Web of Data, developing relevant uses cases and tools for adding “traditional” geospatial data of this *new* scenario to the GIScience community. This provides

evidence that geospatial knowledge graphs are increasing and, hence, entails new ways of thinking about how GIScience and its community should interact with (geospatial) knowledge graphs. In this sense, our approach has allowed us to request retrieve, and consume (geospatial) knowledge graphs based on GeoSPARQL. In this way, this work contributes to filling a gap and makes possible the exploitation of GeoSPARQL-based resources in a distributed environment and through federated (queries) practices, allowing the GIScience community to exploit the richness of the Web of Data.

Concerning the geospatial resources available on the Web of Data, although an increasing number of approaches are adopting GeoSPARQL, we are aware that most major knowledge graphs existing in the Web of Data (e.g., Wikidata or DBpedia) represent geospatial data using the Basic Geo Vocabulary (<https://www.w3.org/2003/01/geo/> (accessed on 23 November 2021)) (WGS84 lat/long). This WGS84 vocabulary is one of the top vocabularies in the context of the Web of Data among the existing vocabularies for describing geospatial knowledge graphs. This basic vocabulary provides a namespace for representing lat (itude), long (itude), and other information about spatially located things, using WGS84 as a reference datum. Considering this scenario, most sources represent features as points and do not support GeoSPARQL. However, an increasing number of initiatives are implementing the mentioned OGC standard. Despite the limitation of our approach to deal with most major geospatial knowledge graphs, we can use more complex geometry types and a comprehensive set of geospatial operators to take advantage of the potential of this geospatial information available on the Web of Data. Likewise, the WGS84 vocabulary support can be included in the future as an extension to collect additional data present in other geospatial knowledge graphs.

Regardless of the shared commitment to knowledge graphs and how, in this line, our work provides an approach for requesting, retrieving, and consuming geospatial knowledge graphs in a federated way, an in-depth knowledge about these resources of the Web of Data is required, since each of these resources can be challenging to use because of the variety of ontologies, ontology structures, and access interfaces.

Our approach is a prototypical implementation developed to show off the GIScience community's potential for the request, retrieval, and consumption of knowledge graphs. In this sense, although our work allows us to deal with complex geometry types based on GeoSPARQL, there are some limitations to handling large volumes of geospatial features with complex geometry because these features are represented with verbose (geometrical) data.

As we mentioned before, we developed our approach on Apache Marmotta, and it was tested through two examples using several (geospatial) knowledge graphs deployed on different triple stores engines like Virtuoso and Apache Marmotta. However, we did not test our approach with other engines, such as Apache Jena, GraphDB, or Parliament, due to the lack of geospatial knowledge graphs deployed on the mentioned triple store engines. Nevertheless, we will work to deploy several geospatial resources regarding some of these engines and perform different tests using some of the recent GeoSPARQL benchmarks [7,47–49].

Regarding the web application development, the current version allows us to display collected resources from federated geospatial queries to the Web of Data as a lightweight alternative to the QGIS connection. However, there are some limitations regarding visualized elements and interactions with geospatial resources. This is due to the current version being an ongoing development, and, therefore, we will develop different releases of this web application.

This article has presented an approach for requesting and retrieving geospatial knowledge graphs through a federated way combining SPARQL 1.1 and GeoSPARQL standards. Moreover, our approach allowed us to consume geospatial knowledge graphs and exploit data directly through federated GeoSPARQL queries, displaying collected data in an application web and through QGIS. Furthermore, our contribution employed geospatial and

non-geospatial knowledge graphs and showed off the potential of our approach through two examples that handle various (geospatial) knowledge graphs from the Web of Data.

Author Contributions: Conceptualization, Luis M. Vilches-Blázquez and Oswaldo Páez; methodology, Luis M. Vilches-Blázquez and Oswaldo Páez; software, Oswaldo Páez and Luis M. Vilches-Blázquez; validation, Oswaldo Páez and Luis M. Vilches-Blázquez; investigation, Oswaldo Páez and Luis M. Vilches-Blázquez; writing—original draft preparation, Luis M. Vilches-Blázquez and Oswaldo Páez; writing—review and editing, Luis M. Vilches-Blázquez and Oswaldo Páez; funding acquisition, Luis M. Vilches-Blázquez. All authors have read and agreed to the published version of the manuscript.

Funding: This work was partially supported by the Instituto Panamericano de Geografía e Historia, Red GeoLIBERO—Consolidación de una red de geomática libre aplicada a las necesidades de Iberoamérica (CYTED program—520RT0010), and Generación de grafos de conocimiento sobre eventos meteorológicos urbanos (SIP-IPN 20210677).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: We kindly thank José Enrique Ortiz Vivar from Universidad de Cuenca (Ecuador) for their technical support during the implementation tasks that allowed us to achieve this work.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Battle, R.; Kolas, D. Enabling the geospatial semantic web with parliament and GeoSPARQL. *Semant. Web* **2012**, *3*, 355–370. [CrossRef]
2. Berners-Lee, T. Linked Data Design Issues. 2006. Available online: <https://www.w3.org/DesignIssues/LinkedData.html> (accessed on 17 November 2021).
3. Heath, T.; Bizer, C. *Linked Data: Evolving the Web into a Global Data Space*; Morgan & Claypool: San Rafael, CA, USA, 2011.
4. Auer, S.; Lehmann, J.; Hellmann, S. *LinkedGeoData—Adding a Spatial Dimension to the Web of Data*; Lecture Notes in Computer Science; Springer: Berlin, Germany, 2009; Volume 5823, pp. 731–746. [CrossRef]
5. Goodwin, J.; Dolbear, C.; Hart, G. Geographical Linked Data: The Administrative Geography of Great Britain on the Semantic Web. *Trans. GIS* **2009**, *12*, 19–30. [CrossRef]
6. Mai, G.; Janowicz, K.; Yan, B.; Scheider, S. Deeply integrating linked data with geographic information systems. *Trans. GIS* **2019**, *23*, 579–600. [CrossRef]
7. Jovanovik, M.; Homburg, T.; Spasić, M. A GeoSPARQL Compliance Benchmark. *ISPRS Int. J. Geo-Inf.* **2021**, *10*, 487. [CrossRef]
8. Kuhn, W.; Kauppinen, T.; Janowicz, K. Linked Data: A paradigm shift for geographic information science. In *Geographic Information Science: GIScience 2014*; Lecture Notes in Computer Science; Duckham, M., Pebesma, E., Stewart, K., Frank, A.U., Eds.; Springer: Cham, Switzerland, 2014; Volume 8728, pp. 173–186. [CrossRef]
9. Perry, M.; Herring, J. (Eds.) *GeoSPARQL—A Geographic Query Language for RDF Data*; OGC®: Rockville, MD, USA, 2012. Available online: https://portal.opengeospatial.org/files/?artifact_id=47664 (accessed on 23 November 2021).
10. Rowland, A.; Folmer, E.; Beek, W. Towards Self-Service GIS—Combining the Best of the Semantic Web and Web GIS. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 753. [CrossRef]
11. Weiss, C.; Karras, P.; Bernstein, A. Hexastore: Sextuple indexing for Semantic Web data management. *Proc. VLDB Endow.* **2008**, *1*, 1008–1019. [CrossRef]
12. Saquicela, V.; Vilches-Blázquez, L.M.; Tello, A. Challenges and trends about smart big geospatial data: A position paper. In *Proceedings of the 2017 IEEE International Conference on Big Data (Big Data)*, Boston, MA, USA, 11–14 December 2017; IEEE: Piscataway, NJ, USA, 2017; pp. 3471–3475. [CrossRef]
13. Usery, E.L.; Varanka, D. Design and development of linked data from the national map. *Semant. Web* **2012**, *3*, 371–384. [CrossRef]
14. Van den Brink, L.; Janssen, P.; Quak, W.; Stoter, J.; Kadaster, J. Linking spatial data: Semi-automated conversion of geo-information models and GML data to RDF. *Int. J. Spat. Data Infrastruct. Res.* **2014**, *9*, 59–85.
15. Vilches-Blázquez, L.M.; Villazón-Terrazas, B.; Corcho, O.; Gómez-Pérez, A. Integrating geographical information in the linked digital Earth. *Int. J. Digit. Earth* **2014**, *7*, 554–575. [CrossRef]
16. Ngonga, A.C. Orchid-reduction-ratio-optimal computation of geospatial distances for link discovery. In *The Semantic Web: ISWC 2013*; Lecture Notes in Computer Science; Alani, H., Kagal, L., Fokoue, A., Groth, P., Biemann, C., Xavier, J., Janowicz, K., Eds.; Springer: Berlin, Germany, 2013; Volume 8218, pp. 395–410. [CrossRef]

17. Vilches-Blázquez, L.M.; Saquicela, V.; Corcho, O. Interlinking geospatial information in the Web of Data. In *Bridging the Geographic Information Sciences; Lecture Notes in Geoinformation & Cartography*; Gensel, J., Josselin, D., Vandenbroucke, D., Eds.; Springer: Berlin, Germany, 2012; pp. 119–139. [\[CrossRef\]](#)
18. Mai, G.; Janowicz, K.; Hu, Y.; McKenzie, G. A Linked Data driven visual interface for the multi-perspective exploration of data across repositories. In *Proceedings of the Fourth International Workshop Visualization and Interaction for Ontologies and Linked Data, CEUR-WS, Kobe, Japan, 17 October 2016; Volume 1704*, pp. 93–101. Available online: <https://www.diva-portal.org/smash/get/diva2:1033953/FULLTEXT02.pdf#page=101> (accessed on 23 November 2021).
19. Sherif, M.A.; Dreßler, K.; Smeros, P.; Ngomoa, A.C.N. Radon: Rapid discovery of topological relations. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence, San Francisco, CA, USA, 4–9 February 2017; AAAI: Menlo Park, CA, USA, 2017*. Available online: <https://www.aaai.org/ocs/index.php/AAAI/AAAI17/paper/viewFile/14199/13759> (accessed on 23 November 2021).
20. Smeros, P.; Koubarakis, M. Discovering spatial and temporal links among RDF data. In *Proceedings of the 9th Workshop on Linked Data on the Web, Montreal, QC, Canada, 11–15 April 2016; ACM: New York, NY, USA, 2016*. Available online: <http://ceur-ws.org/Vol-1593/article-06.pdf> (accessed on 23 November 2021).
21. Jones, J.; Kuhn, W.; Keßler, C.; Scheider, S. Making the web of data available via web feature services. In *Connecting a Digital Europe through Location and Place; Lecture Notes in Geoinformation and Cartography*; Huerta, J., Schade, S., Granell, C., Eds.; Springer: Berlin, Germany, 2014; pp. 341–361. [\[CrossRef\]](#)
22. Wiemann, S.; Bernard, L. Spatial data fusion in spatial data infrastructures using linked data. *Int. J. Geogr. Inf. Sci.* **2016**, *30*, 613–636. [\[CrossRef\]](#)
23. Vilches-Blázquez, L.M.; Saavedra, J. A framework for connecting two interoperability universes: OGC web feature services and linked data. *Trans. GIS* **2019**, *23*, 22–47. [\[CrossRef\]](#)
24. Bishr, Y. Overcoming the semantic and other barriers to GIS interoperability. *Int. J. Geogr. Inf. Sci.* **1998**, *12*, 299–314. [\[CrossRef\]](#)
25. Brachman, R.J. On the epistemological status of semantic networks. In *Associative Networks*; Academic Press: Cambridge, MA, USA, 1979; pp. 3–50.
26. Newell, A. The knowledge level. *Artif. Intell.* **1982**, *18*, 87–127. [\[CrossRef\]](#)
27. McDonald, J.; LevineClark, M. (Eds.) *Encyclopedia of Library and Information Sciences*, 4th ed.; CRC Press: Boca Raton, FL, USA, 2017.
28. Paulheim, H. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semant. Web* **2017**, *8*, 489–508. [\[CrossRef\]](#)
29. Ehrlinger, L.; Wöß, W. Towards a Definition of Knowledge Graphs. *SEMANTiCS* **2016**, *48*, 2. Available online: <http://citeseerx.ist.psu.edu/viewdoc/download> (accessed on 23 November 2021).
30. Herring, J.R. (Ed.) *Implementation Specification for Geographic Information—Simple Feature Access*; Open Geospatial Consortium Inc.: Rockville, MA, USA, 2010. Available online: http://portal.opengeospatial.org/files/?artifact_id=25354 (accessed on 17 November 2021).
31. Schade, S.; Cox, S. Linked data in SDI or how GML is not about trees. In *Proceedings of the 13th AGILE Conference on Geographic Information Science, AGILE, Guimarães, Portugal, 11–14 May 2010*.
32. Stadler, C.; Lehmann, J.; Hoffner, K.; Auer, S. LinkedGeoData: A core for a web of spatial open data. *Semant. Web* **2012**, *3*, 333–354. [\[CrossRef\]](#)
33. Vilches-Blázquez, L.M.; Villazón-Terrazas, B.; Saquicela, V.; de León, A.; Corcho, O.; Gómez-Pérez, A. GeoLinked data and INSPIRE through an application case. In *Proceedings of the 18th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, Online, 2 November 2010; ACM: San Jose, CA, USA, 2010*; pp. 446–449. [\[CrossRef\]](#)
34. Patroumpas, K.; Alexakis, M.; Giannopoulos, G.; Athanasiou, S. TripleGeo: An ETL tool for transforming geospatial data into RDF triples. In *Proceedings of the EDBT/ICDT Workshops, Athens, Greece, 24–28 March 2014; pp. 275–278*. Available online: <https://citeseerx.ist.psu.edu/viewdoc/download> (accessed on 23 November 2021).
35. Kyzirakos, K.; Vlachopoulos, I.; Savva, D.; Manegold, S.; Koubarakis, M. GeoTriples: A tool for publishing geospatial data as RDF graphs using R2RML mappings. In *Proceedings of the 13th International Semantic Web Conference, CEUR, Riva del Garda, Italy, 19–23 October 2014; Volume 1272*, pp. 393–396. Available online: <http://ceur-ws.org/Vol-1401/tc-ssn2014-complete.pdf#page=35> (accessed on 23 November 2021).
36. Brodt, A.; Nicklas, D.; Mitschang, B. Deep integration of spatial query processing into native RDF triple stores. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems, San Jose, CA, USA, 2–5 November 2010; ACM: New York, NY, USA, 2010*; pp. 33–42. [\[CrossRef\]](#)
37. Rowland, A.; Folmer, E.; Beek, W.; Wenneker, R. Interoperability and Integration: An Updated Approach to Linked Data Publication at the Dutch Land Registry. In *Proceedings of the GeoLD 2021 Geospatial Linked Data Workshop 2021—CEUR Workshop Proceedings, Dublin, Ireland, 31 March 2021; Volume 2977*. Available online: <http://ceur-ws.org/Vol-2977/paper3.pdf> (accessed on 23 November 2021).
38. Koubarakis, M.; Kyzirakos, K. Modeling and querying metadata in the semantic sensor web: The model stRDF and the query language stSPARQL. In *Extended Semantic Web Conference, Heraklion, Crete, Greece, 30 May–3 June 2010*; Springer: Berlin/Heidelberg, Germany, 2010; pp. 425–439. [\[CrossRef\]](#)
39. Bereta, K.; Xiao, G.; Koubarakis, M.; Hodrius, M.; Bielski, C.; Zeug, G. Ontop-spatial: Geospatial data integration using GeoSPARQL-to-SQL translation. In *Proceedings of the 15th International Semantic Web Conference, Posters & Demonstrations*

- Track (ISWC), Kobe, Japan, 17–21 October 2016; Volume 1690. Available online: <http://cgi.di.uoa.gr/~koubarak/publications/2016/main.pdf> (accessed on 23 November 2021).
40. Nys, G.A.; Van Ruymbeke, M.; Billen, R. Spatio-temporal reasoning in CIDOC CRM: An hybrid ontology with GeoSPARQL and OWL-Time. In Proceedings of the CEUR Workshop Proceedings, Torino, Italy, 22 October 2018; RWTH Aachen University: Aachen, Germany, 2018; Volume 2230. Available online: <https://orbi.uliege.be/bitstream/2268/228461/1/paper.pdf> (accessed on 23 November 2021).
 41. Vidal, M.E.; Castillo, S.; Acosta, M.; Montoya, G.; Palma, G. On the Selection of SPARQL Endpoints to Efficiently Execute Federated SPARQL Queries. In *Transactions on Large-Scale Data-and Knowledge-Centered Systems XXV*; Lecture Notes in Computer Science; Hameurlain, A., Küng, J., Wagner, R., Eds.; Springer: Berlin/Heidelberg, Germany, 2016; Volume 9620. [CrossRef]
 42. Nikolov, A.; Haase, P.; Trame, J.; Kozlov, A. Ephedra: Efficiently Combining RDF Data and Services Using SPARQL Federation. In *Knowledge Engineering and Semantic Web. KESW 2017*; Communications in Computer and Information Science; Różewski, P., Lange, C., Eds.; Springer: Berlin/Heidelberg, Germany, 2017; Volume 786. [CrossRef]
 43. Hellmund, T.; Schenk, M.; Hertweck, P.; Moßgraber, J. Employing Geospatial Semantics and Semantic Web Technologies in Natural Disaster Management. In Proceedings of the 15th International Conference on Semantic Systems (SEMANTiCS 2019), Karlsruhe, Germany, 9–12 September 2019.
 44. Almobydeen, S.B.; Viqueira, J.R.; Penin, M.L. A Federated Approach for Array and Entity Environmental Linked Data. *XXI Jorn. De Ing. Del Softw. Y Bases De Datos* **2020**, *219*, 385.
 45. Heling, L.M.; Acosta, M. A Framework for Federated SPARQL Query Processing over Heterogeneous Linked Data Fragments. *arXiv* **2021**, arXiv:2102.03269. Available online: <https://arxiv.org/abs/2102.03269> (accessed on 17 November 2021).
 46. Taelman, R.; Van Herwegen, J.; Vander Sande, M.; Verborgh, R. Comunica: A Modular SPARQL Query Engine for the Web. In Proceedings of the Semantic Web—ISWC 2018, Lecture Notes in Computer Science. Monterey, CA, USA, 8–12 October 2018; Vrandečić, D., Bontcheva, K., Suárez-Figueroa, M.C., Presutti, V., Celino, I., Sabou, M., Kaffee, L.-A., Simperl, E., Eds.; Springer: Cham, Germany, 2018; Volume 11137, pp. 239–255. [CrossRef]
 47. Huang, W.; Raza, S.A.; Mirzov, O.; Harrie, L. Assessment and Benchmarking of Spatially Enabled RDF Stores for the Next Generation of Spatial Data Infrastructure. *ISPRS Int. J. Geo-Inf.* **2019**, *8*, 310. [CrossRef]
 48. Albiston, G.L.; Osman, T.; Chen, H. GeoSPARQL-Jena: Implementation and benchmarking of a GeoSPARQL graphstore. *Under Rev. Semant. Web J.* **2019**. *under review*.
 49. Jovanovik, M.; Homburg, T.; Spasić, M. Software for the GeoSPARQL compliance benchmark. *Softw. Impacts* **2021**, *8*, 100071. [CrossRef]
 50. Troumpoukis, A.; Konstantopoulos, S.; Mouchakis, G.; Prokopaki-Kostopoulou, N.; Paris, C.; Bruzzone, L.; Koubarakis, M. GeoFedBench: A Benchmark for Federated GeoSPARQL Query Processors. In ISWC (Demos/Industry). 2020. Available online: <http://ceur-ws.org/Vol-2721/paper558.pdf> (accessed on 23 November 2021).
 51. Tandy, J.; van den Brink, L.; Barnaghi, P. (Eds.) *Spatial Data on the Web Best Practices. W3C Working Group Note*; OGC®: Rockville, MD, USA, 2018; Available online: <https://www.w3.org/TR/sdw-bp/> (accessed on 17 November 2021).
 52. Lehmann, J.; Athanasiou, S.; Both, A.; Rojas, A.G.; Giannopoulos, G.; Hladky, D.; Zaslowski, V. Managing geospatial linked data in the GeoKnow project. In *The Semantic Web in Earth and Space Science: Current Status and Future Directions*; Narock, T., Fox, P., Eds.; IOS Press: Amsterdam, The Netherlands, 2015; Volume 20, pp. 51–78.
 53. Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; Ives, Z. DBpedia: A Nucleus for a Web of Open Data. In *The Semantic Web. ISWC 2007, ASWC 2007*; Lecture Notes in Computer Science; Aberer, K., Choim, K.-S., Noy, N., Allemang, D., Lee, K.-I., Nixon, L., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., et al., Eds.; Springer: Berlin/Heidelberg, Germany, 2007; Volume 4825. [CrossRef]