Geospatial Semantic Web: Architecture of Ontologies

Dave Kolas, John Hebeler, and Mike Dean

BBN Technologies, Arlington, VA USA
{dkolas, jhebeler, mdean}@bbn.com

Abstract. An effective ontology architecture enables the development of a geospatial semantic system that forges multiple geospatial data sources into a powerful cross-discipline knowledge source. This paper suggests types of ontologies that could support a geospatial semantic system. Motivations of each of the ontology types will be expounded, as well as potential areas for standardization by the geospatial community. Finally, the use of this approach within the OGC GSW Interoperability Experiment will be discussed.

1 Introduction

The assimilation of semantic web technologies into the world of geospatial information presents several interesting challenges, of which perhaps one of the most obvious is how to incorporate and leverage existing knowledge representations such as Geography Markup Language [GML], the ISO 19100 series [ISO19100], descriptions of geospatial services, and descriptions of geospatial queries into a coherent set of ontologies. A number of different organizations have begun to convert geospatial knowledge representations, but a more focused approach must be developed and led by the Open Geospatial Consortium[1] if the effort is to succeed. Standardized geospatial ontologies form the foundation upon which many more specific geospatial ontologies should be built. Indeed, this area is one of the primary focuses of the Geospatial Semantic Web Interoperability Experiment [GSW IE] within the OGC, of which BBN is an initiator.

This paper outlines five distinct ontology types in OWL [OWL] that contribute to forming a geospatial semantic system. Each ontology type represents a key role in establishing a rich, dynamic, and flexible geospatial knowledgebase. The five ontology types are:

1. Base geospatial ontology – Provides the core geospatial knowledge vocabulary and knowledge structure.
2. Feature data source ontology – Provides an ontological view of WFS data. Allows WFS and GML data sources to fully participate with knowledgebases and ontologies.
3. Geospatial service ontology – Enables knowledgebase discovery and execution of all registered geospatial services.
4. Geospatial filter ontology – Enables the integration of geospatial relationships into the queries.
5. Domain ontology – Provides a knowledge representation that is organized, customized, and aligned with a specific domain and/or user.

[1] http://www.opengeospatial.org

M.A. Rodríguez et al. (Eds.): GeoS 2005, LNCS 3799, pp. 183 – 194, 2005.

We explore each in turn by using a reference example to highlight each ontology's contributions. In addition, the paper highlights areas that require formal standardization.

2 Example Scenario and Architecture

We construct a simple scenario to illustrate the role and benefits of each ontology type.

2.1 Scenario: An Aircraft Emergency

Imagine a scenario in which a Boeing 747 aircraft is flying over a major metropolitan area. Without warning, the plane begins experiencing engine trouble. Emergency crews scramble to determine the best course of action for the plane. The scramble includes querying many data and knowledge sources. Where are the nearest airports? Are these airports capable of supporting a 747? Would the weather conditions at the airports allow an emergency landing? At which airport will a runway be most easily cleared? What other activities are already underway at each location? Which airport has access to the best medical facilities and mechanical facilities?

Current approaches to this situation require the consultation of many different sources of information, some geospatial and some not. Each would likely be in a proprietary format and require special, independent access methods. Much of the integration of the knowledge would occur manually. The success or failure of the actions would depend upon the emergency crew's inherent knowledge of where to find appropriate data and the timeliness of their decisions. All of this takes time and skill that the flight crew can ill afford.

In the semantic web vision, however, much of this data retrieval and integration could be performed automatically. Ideally, instead of seeking out disparate sources of information, an emergency worker could formulate one query to one system which would break the query down into components, access the appropriate data sources, and return an answer or answers. The gained knowledge is then easily incorporated into the knowledgebase to assist in future, similar situations.

2.2 The Geospatial Ontology Architecture

Fig. 1 shows the proposed ontology architecture in supporting a knowledgebase:

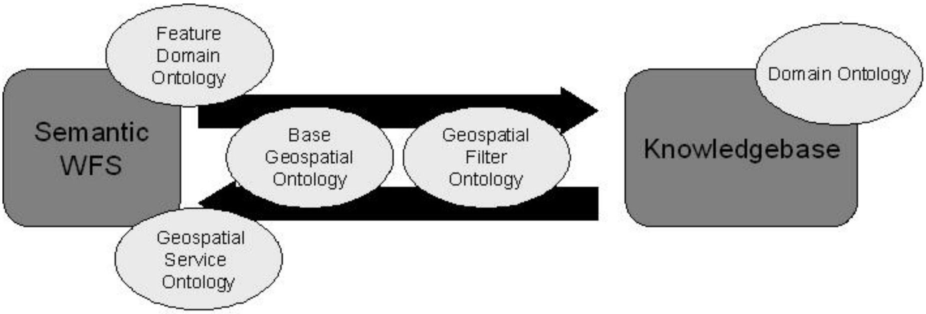

Fig. 1. Geospatial Ontology Architecture

The base geospatial ontology is the core vocabulary that all of the other ontologies must reference. It is a common language and knowledge structure used by both the knowledgebase and the WFS to represent geospatial data, and is analogous to the GML specification [GML]. The domain ontology represents the data from the perspective of a specific user or group. Data is mapped into the domain ontology from the base geospatial ontology and the feature data source. The knowledgebase and the domain ontology are shown on the far right of the diagram to imply that the WFS [WFS] does not require any knowledge of their existence. The filter ontology is a representation of WFS filters as well as geospatial relationships within the query and within the knowledgebase, and is inherently built on base geospatial ontology concepts. It, like the base geospatial ontology, is shared by the WFS and the knowledgebase. The feature data source ontology defines the data from the perspective of the WFS, and is built on base geospatial concepts. While the feature data source ontology is primarily associated with the WFS, the knowledgebase must also have some understanding of it to function as desired. The geospatial service ontology links the base geospatial ontology and the filter ontology to OWL-S [OWL-S] to create semantic web feature service descriptions. Each of these ontologies will now be discussed in detail, with focus given to the base geospatial ontology, the geospatial service ontology, and the geospatial filter ontology which could be standardized by the community.

All five ontologies serve in forming an effective response to the above situation. A query built with the domain ontology initiates a request on the worker's behalf and from the worker's perspective. In order to express the geospatial relationships and filter the data geospatially, the query uses the geospatial filter ontology and the base geospatial ontology. Parts of the query can then be translated for the WFS, which is located via the geospatial service ontology. The service ontology describes the types of services offered by the WFS. The WFS data is returned using its associated ontology, the feature data source ontology. The feature data source ontology translates and maps the underlying formats of WFS into an ontology. This bridges the technology gap between the WFS and ontologies. The WFS responds to the request by forming a response in the feature data source ontology. This data is then mapped to the domain ontology and formatted for the user.

3 Base Geospatial Ontology

The base geospatial ontology forms the ontological foundation of geospatial information. It provides a common base to which geospatial knowledge representations can be linked.

3.1 Motivations for a Common Geospatial Ontology

The motivations for creating a base geospatial ontology start with the same goals that led to the creation of GML:

- A standard way to communicate geospatial data between applications
- An expandable definition of types such that applications can extend the types with their own data
- The capacity to link geospatial application schemas

The current XML representation of GML can realize all of these goals, but a base geospatial ontology extends its power with the significantly greater expressiveness of OWL and the ability to link this data to knowledge outside the geospatial realm. This expands the overall usefulness of the geospatial data while enriching it with complementary information.

3.2 Relationship to GML

The base geospatial ontology should leverage the work done in creating GML by including the same geometry types that are in GML. By using the same concepts that already exist, we enable easier translation from GML to the ontology. In fact, an initial version of this ontology might simply be a conversion of the schemas into OWL. While this would not add any expressiveness per se, it would allow the other ontologies to link to the geospatial content in a widely used manner. As development on the ontology continues, further refinement of the relationships between the types could be added to fully utilize the descriptive power of OWL.

3.3 Linking to the Base Geospatial Ontology

Each of the other ontologies to be discussed has a direct relationship to the base geospatial ontology. Both the domain ontology and the feature data source ontology could extend the types within the base geospatial ontology through OWL just as a WFS can extend the base GML types through XML schema. The other two ontologies, the filter ontology and the service ontology, will necessarily contain properties which use base geospatial types.

3.4 Opportunities for Development and Standardization

Because the base geospatial ontology is the foundation upon which the other ontologies must be built, it should be the primary focus of standardization among the geospatial community. The substantial work that has gone into the creation of the GML schemas as well as the ISO 19100 series of standards will be of enormous benefit both in scoping such an ontology and defining its content.

4 Geospatial Service Ontology

The geospatial service ontology conforms to the OWL-S specification enabling full knowledgebase queries against service offerings. This enhances and extends the current offerings via web services.

4.1 Motivations for Semantic Service Descriptions

While some of the data services that the knowledgebase employs may be known a priori, an effective solution requires the ability to discover and utilize these services automatically. In this way, the goals for the interaction are closely aligned with those of OWL-S [OWL-S]: automatic discovery, invocation, and composition of semantic web services.

In order for a service to be discovered automatically, it must advertise in an abstract sense what it is and what it can do. This corresponds to the concept in OWL-S of a *profile*. The profile of an OWL-S description describes who provides the service, what the service does, and other properties characterizing the service. This description allows the knowledgebase to decide whether or not a particular service is appropriate to its needs, particularly, whether or not the service has data relevant to the query being asked of the knowledgebase.

Automatic service invocation requires both an abstract concept of how the individual services work and the concrete way in which the knowledgebase must interact with them. OWL-S provides these details through the *process model* and *grounding*. For each individual service provided, the process model states the *inputs, outputs, preconditions, and effects.* This means for each service, our knowledgebase knows what it needs to send to the service, what it will get back, what the state of the world must be for the service to succeed, and what the state of the world will be when it finishes. This allows our knowledgebase to ensure that the service meets its needs, as well as knowing what data it must provide in order to achieve the desired result. The grounding brings the interaction down to the most concrete level, telling the knowledgebase which ports, protocols, and encodings to use for invocation.

Finally, the semantic services description enables our knowledgebase to create composite services. While this may initially seem irrelevant to fetching data from a Web Feature Service, there are situations in which our knowledgebase could certainly profit from its usage. Consider a Web Feature Service with a semantic description that does not have the ability to return its feature data in OWL. The knowledgebase could then chain this service to one that can translate GML data of this WFS's particular type schema to OWL of a given ontology.

4.2 Linking OWL-S Service Descriptions and the WFS Specification

It is easy to draw parallels from the parts of an OWL-S service description to the various parts of the Web Feature Service specification [WFS] and a Web Feature Service's typical advertisement, the capabilities document. However, when drawing the parallels in the opposite direction, there are concepts that are unique to the interaction with a Web Feature Service and have not yet been addressed.

In order for the knowledgebase to make use of the WFS, it needs to not only know what the WFS does and how to interact with it, but also what types of content the WFS can return. This content is defined both in terms of feature types and the extent of the data available within those types. Not only does a WFS tell you that it can find airports, for example, but also that it can find airports between a certain set of latitudes.

The scope of OWL-S, however, does not by default have a place for such statements. To add this content to the service description, an ontological definition of the remaining concepts is necessary. This leads to the definition of the Geospatial Service Ontology.

4.3 Usage of the Geospatial Service Ontology

Because the Geospatial Service Ontology would allow the service description to indicate the type of content that can be found on the server and thus allow the knowledgebase to decide if the service is appropriate, it seems most likely to fit into the realm of the profile. Along with the listing of the types of features that the server can return, the service description would probably need to provide a reference to the ontology that defines these types. This function would be analogous to the *describeFeatureType* operation presented by current Web Feature Services, allowing the knowledgebase to find the ontology from which the types are derived.

It is likely that the service descriptions of various Web Feature Services would be quite similar. In fact, much of the process model and grounding could be nearly identical for different Web Feature Services. This would allow significant reuse from WFS to WFS, and thus quicker development, as well as providing potential grounds for standardization.

4.4 The Relationship Between Semantic Service Descriptions and Registries

It is important to note that semantic web service descriptions are complementary to, and not a replacement for, service registries. By providing information about what the WFS offers, the registry then only needs to catalog these descriptions and provide an interface to search them. Exactly how this interface might work is a nontrivial topic that is outside the scope of this paper.

Useful for discussion, however, is how a registry or catalog might fit into the architecture. While the initial description of the architecture was kept as simple as possible, a likely scenario, especially long-term, would add interaction with a service registry before the interaction with the Web Feature Service. In this situation, before the knowledgebase made any contact with a WFS, it could contact the registry with the type of data it requires and get back descriptions of available appropriate WFS's. Another advantage of the semantic service description for a WFS is that it could participate in both specialized WFS registries as well as more general semantic web service registries. This could allow geospatially inclined semantic applications to make the most efficient use of the available servers, while still allowing more general semantic applications to interact with the WFS.

4.5 Opportunities for Development and Standardization

Since the specification for Web Feature Services falls under the purview of the Open Geospatial Consortium, it seems like a logical progression for the definition of the Geospatial Service Ontology and potentially a standard piece of an OWL-S service description for a WFS to be standardized by the OGC as well. Some of the knowledge of which components standardization would benefit will hopefully be derived from the OGC's Geospatial Semantic Web Interoperability Experiment, which is currently in progress.

5 Geospatial Filter Ontology

The Open Geospatial Consortium currently defines a filter encoding [Filter Spec] for use with Web Feature Services. This encoding is used to represent some geospatial and logical relationships that can be used to filter results of a *getFeature* request. In our scenario, we could construct a filter that would represent all features within a 100 mile radius of the plane's current location:

```
<Filter>
    <DWithin>
        <PropertyName>Geometry</PropertyName>
        <gml:Point>
            <gml:coordinates>(plane's coordinates)</gml:coordinates>
        </gml:Point>
        <Distance units='miles'>100</Distance>
    </DWithin>
</Filter>
```

5.1 A Common Language for Geospatial Relationships

While the filter encoding is very effective for interaction with a Web Feature Service, it could be rendered significantly more powerful if built into an ontology. The combination of the concepts of this filter ontology and the concepts the base geospatial ontology would result in a standardizable language that could be used to describe all spatial relationships in a way that is widely accepted and understood, thus useful for the expression of geospatial knowledge outside the realm of the Web Feature Service.

As geospatial semantic web technology becomes more commonplace, there will be an increasing need for a common language for geospatial relationships. For example, the word “within” could conjure many meanings to many people. However, if the concept of “within” is defined in a spatial relationship ontology, then its meaning can be interpreted without ambiguity through an end to end semantic system such as the one described here. There are four areas within the system where the benefits of this filter ontology are immediately apparent: the semantic description of the Web Feature Service; the interaction with the Web Feature Service’s *getFeature* operation, the creation of the client-side query and its representation, and the creation of semantic rules, rule functions, and their implementations within the knowledgebase. Each of these uses will be addressed individually, starting with those related to the Web Feature Service directly.

5.2 The Filter Ontology and the Semantic Web Feature Service

The use of the filter ontology as related to the semantic services description of the WFS is reasonably straightforward. The filter ontology gives the service description a way to reference which types of filters it does and does not support. This could be given by referencing the filter types from the filter ontology that are supported.

The next use of the filter ontology is in performing the actual *getFeature* operation. In a truly semantic system, this request could be formatted in OWL, allowing for easy integration with semantic clients like our knowledgebase. The existence of the filter ontology enables the building of a filter for a feature request as OWL instance data.

```
<filter:DWithin rdf:ID=featuresWithinRadius">
    <filter:propertyName rdf:resource="&gml;#location" />
    <filter:measuredFromGeometry
            rdf:resource="#CurrentAircraftCoordinates" />
    <filter:distance>
```

```
            <units:Miles>
                <units:value>100</units:value>
            </units:Miles>
        </filter:distance>
    </filter:DWithin>
    <gml:Point rdf:ID="CurrentAircraftCoordinates">
        …
    </gml:Point>
```

5.3 The Filter Ontology and the Knowledgebase

Note that to this point no additional power of the filter ontology has been demonstrated, but existing capabilities have been mapped to its use. The additional power derived from using the filter ontology in the semantic web system comes from its use in the client side of the architecture. The first such use is utilizing the ontology to formulate geospatial semantic queries.

Semantic queries with geospatial relationships initially present a problem for RDF query processes; some of the relationships, DWithin and Beyond, are n-ary. This clearly prevents them from being expressed as one simple RDF property. One potential solution to this problem involves user defined functions within the query. This approach is supported by the emerging query language SPARQL [SPARQL]. While this method will certainly enable the geospatial relationships to be placed in the query, it has a few significant disadvantages to using the filter ontology to represent them. First, it requires whatever front-end client that is being used to be able to specially format its semantic queries in this way. While this may not be a significant problem for systems which are focused specifically on geospatial information processing, it makes it less likely that a system designed to build queries from existing ontological concepts will be able to build specialized geospatial queries. Moreover, the addition of more spatial relationships then would require enhancing both the clients and the knowledgebase instead of just the knowledgebase. Secondly, using the special functions within the query could place the burden of calculation on the query processor instead of the knowledgebase. Since the trend now is for integrated rule processing at the storage layer, it seems beneficial to use the mechanisms that will likely already be in place rather than add an additional dependency.

The final area in which the filter ontology could be used is in rule processing within the knowledgebase. While queries that are handled entirely by a WFS could be filtered entirely by the WFS, those that incorporate existing data or require the processing of combinations of data will require the spatial calculations to be performed within the knowledgebase. This ability could be manifested in a set of SWRL built-ins [SWRL], implementations of these built-ins, and SWRL rules making use of the built-ins.

By referencing the same spatial relationships here as are used in the query, there is no requirement for any special knowledge of geospatial relationships in the client front-end; that is, no knowledge about how to process them is required above and beyond how to process any semantic relationship.

5.4 Opportunities for Development and Standardization

Just as the filter encoding specification is standardized by the OGC, it seems logical for such a standards body to develop and approve a specification for an ontology based thereupon. Aside from the standardization of this ontology, it could be beneficial for the OGC to participate in the development of the geospatial SWRL built-ins and potentially a standard set of SWRL rules to imply the geospatial relations (e.g. propagation of containment relations).

6 Domain Ontology

The domain ontology is not a specific ontology, but rather a class of ontologies representing the perspective of the user community that requires geospatial information. As such, it could be made up of or derived from many other ontologies. The domain ontology will very likely be built upon many public ontologies, and could very well be public itself; however, it will be considered private in this scenario because only the client side of the architecture requires any knowledge thereof.

The domain ontology represents the data in the knowledgebase from the perspective of subject matter experts in the relevant domain. This ontology relates the geospatial features and all other data into terms that the user understands. For example, in our emergency scenario, an airplane is something that flies between airports; from the domain of airplane manufacturing, however, the same airplane could be viewed as a product.

The primary purpose of the domain ontology is to represent the concepts over which the user will query. In fact, this is the standard upon which a good domain ontology should be measured; if the user has the vocabulary to ask the questions that they want to ask, then the domain ontology is successful.

In order for the knowledgebase to successfully interact with the Web Feature Service, the domain ontology must have a connection to the base geospatial ontology. This link could be achieved in a number of ways, including subclassing of base geospatial ontology concepts or mapping from base concepts to domain concepts through rules. Either way, this link provides the means to use the previously mentioned geospatial relationships of the filter ontology within queries on the domain ontology.

Due to the domain-specific nature of these ontologies, it is unlikely that they would be general enough with respect to geospatial concepts to be targeted for standardization by the geospatial semantics community.

7 Feature Data Source Ontology

Like the domain ontology, the feature data source ontology is a class of ontologies. However, this class of ontologies is used to represent the domain of the Web Feature Service as opposed to that of the client knowledgebase. Perhaps even more than the domain ontology, the feature data source ontology will almost certainly be built upon or derived from ontologies or data models in common use.

The feature data source ontology is quite analogous to the feature type schemas returned by the WFS *describeFeatureType* operation. It serves exactly the same purpose: to describe the type of data that will be returned beyond the base GML, or in this case base geospatial ontology, types.

The separation of feature data source and domain ontologies also accommodates semantic integration among heterogeneous data sources. For example, WFSs containing aeronautical features corresponding to the US DAFIF [DAFIF] and European AIXM [AIXM] standards could be mapped into a common aeronautical domain ontology.

Clearly, for this additional data to be of use to the client system there must exist a mapping from the feature data source ontology to the client's domain ontology. However, it is not necessary for a client which understands the base geospatial ontology to understand everything or anything from this ontology; the data is simply there if it is desired.

In the long term, it is expected that some ontologies that represent current data schemas as well as new types of data will be developed and become publicly available. This is necessary for semantic clients to make full use of such semantically enabled Web Feature Services. These ontologies are more likely to be standardized by particular communities of interest than by the geospatial community as a whole.

8 The Ontologies in Action: GSW Interoperability Experiment

This ontology approach is currently being evaluated as part of the OGC GSW Interoperability Experiment. Rather than attempting to create full versions of any of these ontologies, we will attempt to validate the approach by creating minimal parts of each ontology to be used as a thread through the system. The current working scenario is very similar to the one described earlier; a specific type of aircraft needs to find a suitable place to land near a city. We have created a knowledgebase which holds OWL data and receives queries in the domain ontology. These queries are then decomposed and translated into appropriate WFS queries. Data returned from the WFS is translated from GML to OWL, and SWRL business rules are applied. The result is then returned to the user.

8.1 Domain Ontology

The domain ontology in this example is that of airports. The ontology defines airports, runways, and other physical features of airports, as well as what it means for an airport to be able to support the plane. It currently also contains the concept of a City, within a different sub-ontology of what would be considered the domain ontology. The domain ontology links to the base geospatial ontology in that a city has Point geometry data.

8.2 Feature Data Source Ontology

The feature data source ontology currently being used is based on the DAFIF schema. This allows for representation of data served by a WFS in DAFIF format. Just as the DAFIF schema links to GML constructs, the DAFIF ontology links to the base geospatial ontology.

8.3 Geospatial Filter Ontology

The filter ontology is designed to mimic the OGC Filter Encoding specification, for easy translation into filters that can be placed in WFS queries. Currently the filter DWithin is being used in the query to define the area in which the airports are being requested. Naturally, this ontology references the base geospatial ontology where the filters require a geometry.

8.4 Base Geospatial Ontology

Currently the base geospatial ontology has only one type of geometry in it: a Point. The point is clearly the simplest form of geometry that we could use, but it nonetheless demonstrates how each of the other ontologies link to base geospatial ontology constructs.

8.5 Geospatial Service Ontology

This is the only ontology concept that has not yet been created or put to use. Using semantic constructs for geospatial service discovery is a current area of research for the Interoperability Experiment.

8.6 Current Usage

A query is formulated for the system using primarily the domain ontology and using the geospatial filter ontology to define the spatial relationships. This query is processed by the knowledgebase, and an appropriate query to what is now a known WFS is formed. The WFS call is made, and the resulting GML is processed into OWL in the form of the DAFIF feature data source ontology. This data is then mapped into the domain ontology through SWRL rules, and the query results are constructed from the product.

9 Conclusions

Developing a rich, standardized set of geospatial ontologies will significantly advance the usefulness and effectiveness of geospatial data regardless of format. This requires the involvement of the whole community to both accurately incorporate the significant progress-to-date and enable the full benefits of the semantic web. Our approach suggests the use of five distinct ontology types to produce a flexible, powerful semantic system – three of which require some form of standardization. Developing a base geospatial ontology, a geospatial service ontology, and a filter ontology to represent geospatial relationships should be goals of the community to further semantic interoperability. The standardization of such ontologies by the OGC would give significant momentum to efforts attempting to create end-to-end semantic geospatial systems, as well as those related to integration of geospatial and non-geospatial data sources.

References

[AIXM] Aeronautical Information Exchange Model. http://www.eurocontrol.int/ais/aixm/

[GML] Cox, S., Daisey, P., Lake, R., Portele, C., Whiteside, A., eds. OpenGIS Geography Markup Language (GML) Implementation Specification, Version 3.0. *Open Geospatial Consortium*, 01/29/2003. https://portal.opengeospatial.org/files/?artifact_id=7174,2003.

[DAFIF] Digital Aeronautical Flight Information File. https://164.214.2.62/products/digitalaero/index.cfm.

[SWRL] Horrocks, I., Patel-Schneider, P.F., Boley, H., Tabet, S., Grosof, B., Dean, M. SWRL: A Semantic Web Rule Language Combining OWL and RuleML. *W3C Member Submission,* http://www.w3.org/Submission/2004/SUBM-SWRL-20040521/,05/21/2004.

[ISO 19100] ISO/TC 211 Geographic Information/Geomatics. http://www.isotc211.org/scope.htm.

[GSW IE] Lieberman, J., et al. Geospatial Semantic Web Interoperability Experiment. Open Geospatial Consortium. http://www.opengeospatial.org/initiatives/?iid=168.

[OWL-S] Martin, D., et al. OWL-S: Semantic Markup for Web Services. *W3C Member Submission*. http://www.w3.org/Submission/2004/SUBM-OWL-S-20041122/, 11/22/2004.

[OWL] McGuinness, D.L., van Harmelen, F., eds. OWL Web Ontology Language Overview. *W3C Recommendation*. http://www.w3.org/TR/2004/REC-owl-features-20040210/, 02/10/2004.

[SPARQL] Prud'hommeaux, E., Seaborne, A., eds. SPARQL Query Language for RDF. *W3C Working Draft*. http://www.w3.org/TR/2005/WD-rdf-sparql-query-20050419/, 04/19/2005.

[Filter Encoding] Vretanos, P., ed. Filter Encoding Implementation Specification, *Version 1.1. Open Geospatial Consortium.* http://portal.opengeospatial.org/files/?artifact_id=8340, 05/03/2005.

[WFS] Vretanos, P., ed. Web Feature Service Implementation Specification, *Version 1.1. Open Geospatial Consortium.* https://portal.opengeospatial.org/files/?artifact_id=8339, 05/03/2005