POLITECNICO DI TORINO
Repository ISTITUZIONALE

"Ontological Representation of Constraints for Geographical Reasoning"

Original
"Ontological Representation of Constraints for Geographical Reasoning" / Torta, G., Ardissono, L., Corona, M., La Riccia, L., Voghera, A.. - ELETTRONICO. - 2:(2018), pp. 136-147. (10th Int. JointConf. on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2018) Seville, Spain 18-20/09/2018) [10.5220/0006960901360147].

Availability:
This version is available at: 11583/2715326 since: 2018-10-18T10:49:11Z

Publisher:
INSTICC

Published
DOI:10.5220/0006960901360147

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

12 September 2026

Ontological Representation of Constraints for Geographical Reasoning

Gianluca Torta[1], Liliana Ardissono[1], Marco Corona[1], Luigi La Riccia[2] and Angioletta Voghera[2]
[1]*Dipartimento di Informatica, Università di Torino, Torino, Italy*
[2]*Dipartimento Interateneo di Scienze, Progetto e Politiche del Territorio, Torino, Italy*
{*gianluca.torta, liliana.ardissono*}*@unito.it, marco.corona@edu.unito.it,* {*luigi.lariccia, angioletta.voghera*}*@polito.it*

Keywords: Geographic Knowledge, Geographical Constraints, GeoSPARQL, Ecological Networks, Urban Planning.

Abstract: We describe a framework that supports multiple types of constraint-based reasoning tasks on a geographic domain, by exploiting a semantic representation of the domain itself and of its constraints. Our approach is based on an abstract graph representation of a geographical area and of its relevant properties, for performing the reasoning tasks.
As a test-bed, we consider the domain of Ecological Networks (ENs), which describe the structure of existing real ecosystems and help planning their expansion, conservation and improvement by introducing constraints on land use. While some previous work has been done about supporting the verification of compliance of fully specified ENs, we aim at taking a significant step further, by addressing the automatic suggestion of suitable aggregations of land patches into elements of the EN. This automated generation of EN elements is relevant to support the human planner in the design of public policies for land use because it leverages automated tools to carry out a possibly lengthy and error-prone task.

1 INTRODUCTION

This paper describes a framework supporting constraint-based reasoning tasks on a geographic domain. We take the motivating example of Ecological Networks (ENs) (Bennett and Mulongoy, 2006) as a starting point to define a more general geographic reasoning model in which both the domain and its constraints are uniformly represented as OWL ontologies (W3C, 2017). In this way, we can benefit from the expressiveness provided by standard Semantic Web languages in the specification of domain knowledge.

ENs have been introduced to describe the structure of existing real ecosystems and help planning their expansion, conservation and improvement by imposing constraints on land use. However, up to now, the design of the ENs structure, and the definition of public policies proposing land use transformations to comply with the EN constraints, have been carried out manually, in a lengthy and possibly error-prone type of activity. This is due to the informal specification of ENs, which consists of a rather large number of guidelines expressed in Natural Language.

In this paper, we propose to represent both the EN domain and its constraints in a semantic knowledge representation language, with the aim of supporting a number of automated reasoning tasks concerning the design of the structure of the EN, as well as urban and regional plans and projects for transforming a geographic area so that it (better) complies with the EN specification. A desired side effect of those activities is to construct the social awareness on bindings and opportunities related to ENs for quality of life.

Specifically, we aim at developing a framework for the automated suggestion of suitable aggregations of land patches of a geographic area into elements of an EN. The main contributions of this paper are:

- An ontological representation of ENs and of the related constraints on land use which supports knowledge sharing and semantic reasoning. Our work fully adheres to the GeoSpatial Semantic Web paradigm (Janowicz et al., 2012a), which promotes standard knowledge representation languages to maximize data and application interoperability.
- A model for the automated generation of a graph-based, abstract representation of a geographical area, which specifies the ecological and land use characteristics of land patches, and their adjacency relations. Specifically, the nodes of the graph represent homogeneous pieces of land, each one associated with a given set of defined properties; e.g., naturalness, or irreversibility of the land use. The arcs represent abstractions of the adja-

Torta, G., Ardissono, L., Corona, M., Riccia, L. and Voghera, A.
Ontological Representation of Constraints for Geographical Reasoning.
In *Proceedings of the 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2018) - Volume 2: KEOD*, pages 136-147
ISBN: 978-989-758-330-8
Copyright © 2018 by SCITEPRESS – Science and Technology Publications, Lda. All rights reserved

cency relations between nodes.

- An extensible framework that exploits the graph-based model for performing reasoning tasks on the EN domain, guided by the knowledge encoded in both the domain and constraints ontologies. As a proof-of-concept, the current implementation of our framework is equipped with two reasoners optimized to solve specific types of tasks: i.e., finding paths that connect certain areas, and clustering areas characterized by the same given properties.

Even though the ultimate goal of our work is to provide ontologies and specialized reasoners that allow a fully-automated, mixed-initiative management of reasoning tasks on geographical areas, in this paper, we focus on tasks relevant to the ENs domain.

The remainder of this paper is organized as follows: Section 2 provides some background on semantic knowledge representation and ENs. Sections 3 and 4 present our knowledge representation and reasoning model. Section 5 describes the framework implementation and Section 6 positions it in the related literature. Sections 7 and 8 present our future work and conclude the paper, respectively.

2 BACKGROUND

2.1 Semantic Representation of Geographic Information

The GeoSpatial Semantic Web vision, presented in (Janowicz et al., 2012a), advocates for representing geographical information by means of ontologies suitable for explicitly describing concepts and relations among concepts. This is important to support a conceptual notion of data interoperability, which goes beyond the adoption of a common representation format, and is aimed at enabling correct data interpretation and inferences in geographical reasoning.

Even before the birth of this vision, several geographical ontologies were developed to support information sharing and retrieval. For instance, Fonseca et al. proposed an ontology to classify geographic elements with respect to geometric characteristics and attribute values; i.e., semantic features (Fonseca et al., 2002a). Moreover, they discussed semantic granularity, i.e., the level of detail at which geographic objects are described, focusing on interoperability issues; see (Fonseca et al., 2002b).

Some general ontologies have been published to share toponyms and generic geographic concepts; e.g., the GeoNames Mappings ontology (GeoNames.org, 2018) based on the GeoNames database (http://www.geonames.org/). Researchers have also focused on spatial granularity to describe toponyms at different levels of detail; e.g., see (Palacio et al., 2015). Other ontologies have been developed to provide a semantics of Volunteer Geographical Information. For instance, LinkedGeoData links crowsourced OpenStreetMap (OSM) information to GeoNames and others ontologies (Janowicz et al., 2012b). Moreover, focusing on Open Data and crowdsourced data, Baglatzi et al. (Baglatzi et al., 2012) bridged the gap between ontological and crowdsourcing practices via ontological alignment. Furthermore, OSMonto (Codescu et al., 2011) has been defined to structure the implicit ontology of OpenStreetMap (OpenStreetMap Contributors, 2017) tags and to format it into an RDF (W3C, 2017b) scheme.

Besides the work supporting geographic information sharing and retrieval, more specific ontologies have been defined to describe fine-grained aspects of geographical objects. Of primary importance for the present work is the GeoSPARQL ontology (OGC, 2012), which describes geographical objects supporting the specification of their geometry, as well as topologic relations among different objects.

Other ontologies relevant to the present work are the ones that model the types of *land use/cover*, e.g., LBCS-OWL2 (Montenegro et al., 2011), HarmonISA (Hall and Mandl, 2006). As we shall see, we have currently adopted a taxonomy based on the Land Cover Piemonte (LCP) cartography (Provincia di Torino, 2014), mainly because the experimental data available to us was tagged according to it, see project (Città Metropolitana di Torino, 2014).

2.2 Ecological Networks

Ecological Networks (ENs) have been proposed to preserve biodiversity by reducing the process of nature fragmentation caused by the development of new urbanizations, infrastructural networks and intensive agriculture (Jongman, 1995). Recently, there has been an exponential growth of urban land use towards more natural spaces: external urban areas (uncultivated or abandoned cultivated land, burnt areas, degraded forests) have often been confined from urban and regional planning to an "inessential" position and sometimes simply considered as "waiting areas for a new urbanization". As reported in (Bennett and Mulongoy, 2006), "ecological networks share two generic goals, namely:

1. Maintaining the functioning of ecosystems as a means of facilitating the conservation of species

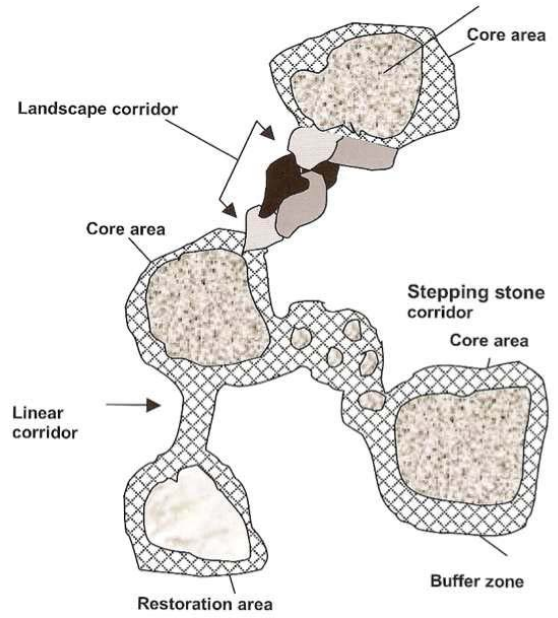

Figure 1: Ecological Network representation, from (Bennett and Mulongoy, 2006).

and habitats.

2. Promoting the sustainable use of natural resources in order to reduce the impacts of new urbanizations on biodiversity and/or to increase the biodiversity value of managed landscapes".

Despite the Protected Areas and *Natura 2000* sites are now considered the backbone of European policy for biodiversity, the increasing expansion of urbanizations and infrastructural networks is challenging the conservation of natural habitats for the preservation of animal species and plant varieties. The policies for the improvement of EN are necessary to overcome the fragmentation of the habitats and natural areas, which is the main cause of biodiversity loss in Europe. The consequences of these processes can be summarized in the following phenomena (Benedict and McMahon, 2002):

1. The substantial loss of natural areas due to urban development.
2. The fragmentation of natural areas into smaller, disconnected patches.
3. The degradation of wetlands, which have an important ecological function for the control of water flows, the ability to block sediments, for the support of plant and animal species, etc. (stepping stone function) and for the ability to provide nutrients for the ecosystems.
4. The inability of ecosystems to respond to changes and find a new ecological balance: that is to say a significantly reduced resilience.
5. The loss of ecosystem services, such as the control of water and the filter functions for pollutants.
6. The increased economic costs for public services, due to the response to natural disasters deriving from human footprint.

An Ecological Network can be defined as an interconnected system consisting of territorial areas that include natural and semi-natural habitats. The typical representation of an EN is a network of core areas interconnected by corridors; see Figure 1. According to (Bennett and Mulongoy, 2006):

- *Core Areas* are the areas "where the conservation of biodiversity takes primary importance, even if the area is not legally protected".
- *Buffer Zones* are the areas which "protect the network from potentially damaging external influences and which are essentially transitional areas characterized by compatible land uses". They are important to safeguard and increase the stability of the core areas.
- *Corridors* "serve to maintain vital ecological or environmental connections by maintaining physical (though not necessarily linear) linkages between the core areas".
- *Sustainable-use areas* are zones "where opportunities are exploited within the landscape mozaic for the sustainable use of natural resources together with maintenance of most ecosystem services".

So far, the work about ENs has mostly focused on the following perspectives:

1. Ecological network analysis has dealt with studying and simulating in a mathematical way the interaction between organisms within the ecosystem, the dynamics of the relations among species, the existence of dynamical bottlenecks in the functioning of the ecosystems, etc.; see, e.g., (Fath et al., 2007; Ulanowicz, 2004; Fath et al., 2007; Lurgi and Robertson, 2011; Gobluski et al., 2016; Pilosof et al., 2017). This research line has a different goal than ours. In fact, it aims at predicting the evolution of an EN starting from an initial status, but it does not deal with representing and reasoning with the land use constraints.
2. Several ENs have been implemented at different scales (from European down to municipality), resulting in the production of different guidelines on land use, and planning documents, expressed in Natural Language; e.g., (Council of Europe, 2000; Bennett and Wit, 2001; Bennett and Mulongoy, 2006). However, their linguistic specifi-

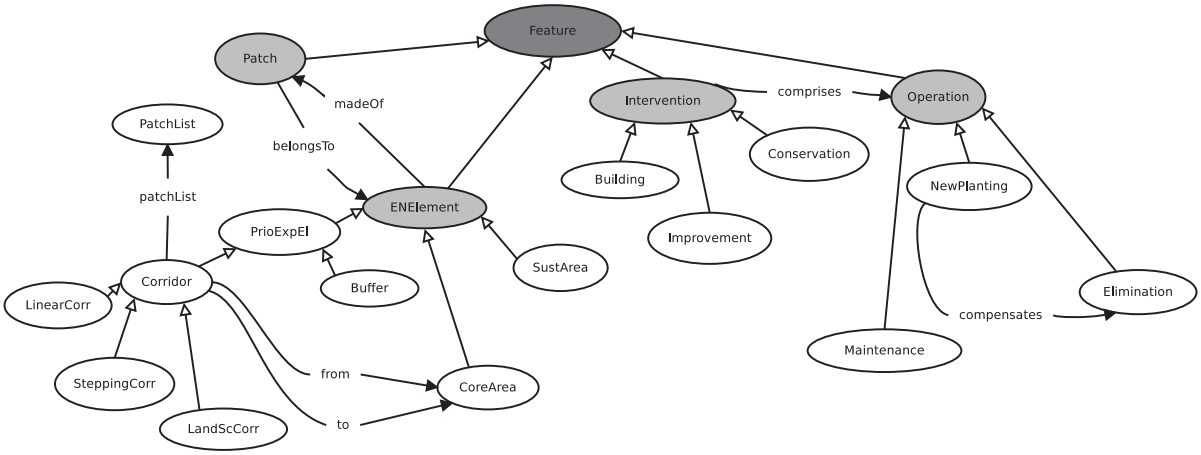

Figure 2: A portion of the EN Ontology.

cation makes it hard to present EN elements in geographical maps, which have to be manually created. Moreover, it challenges an automated verification of EN constraints.

To the best of our knowledge, the only work pursuing the automation of a task needed for EN project planning is the one described in (Torta et al., 2017; Torta et al., 2018), which focuses on EN validation. As discussed later on, we extend that work by adopting a unified semantic approach for a holistic representation of knowledge about ENs and about their constraints. Moreover, we adopt a more general approach to the supported reasoning tasks.

3 KNOWLEDGE REPRESENTATION

3.1 Ecological Networks Representation

Our EN Ontology describes the main concepts and relations of Ecological Networks. We defined it starting from the specifications produced in project "Experimental activity of participatory elaboration of ecological network" (Città Metropolitana di Torino, 2014), conducted by the Metropolitan City of Turin (Italy)[1] in collaboration with Polytechnic of Turin[2] and ENEA[3]. Specifically, in the present work we extend and refine the ontology proposed in (Torta et al., 2017; Torta et al., 2018). The above mentioned project aimed at defining a proposal for the Ecological Network implementation at the local level in two pilot municipalities near Turin. The proposed approach was aimed at guiding local Public Administrations with measures to limit anthropogenic land use and, where possible, orient and qualify the conservation of ecosystem services.

Figure 2 shows the main classes of the EN Ontology, which is developed in OWL (W3C, 2017), and is augmented by the GeoSPARQL ontology (OGC, 2012). The graphic notation is borrowed from (Van Hage et al., 2011); in particular, arrows with open arrow heads symbolize subclass relationships between the classes, while regular arrows connect domains and ranges of properties [4]. The top class of the EN Ontology is the *Feature*, a class defined in the GeoSPARQL ontology and colored in dark gray for easy recognition. A *Feature* has a *Geometry* on the 2D plane, and can thus be used to represent points, lines, and areas on a geographic map. GeoSPARQL defines a number of topological geometric relations between *Feature*s that correspond to fundamental relations (such as *intersects*, equals, contains, etc.) known in the literature as Simple Features (Open Geospatial Consortium et al., 2011).

In order to define ENs, we specify the *Feature* class into the four hierarchies of classes representing the core of the domain. The roots of these hierarchies are colored in gray.

- *ENElement* represents an element of the EN. This is either a Core Area (*CoreArea*), a Sustainable-use area (*SustArea*), or a Priority Expansion Element, which is further specialized to a Corridor (*Corridor*) or Buffer Zone (*Buffer*).
- *Intervention* represents an intervention for building, improving or conserving the EN.
- *Operation* represents a specific operation of elimination (*Elimination*), construction (*NewPlanting*),

[1] www.cittametropolitana.torino.it/
[2] www.polito.it
[3] www.enea.it
[4] The diagrams were created with the Dia tool.

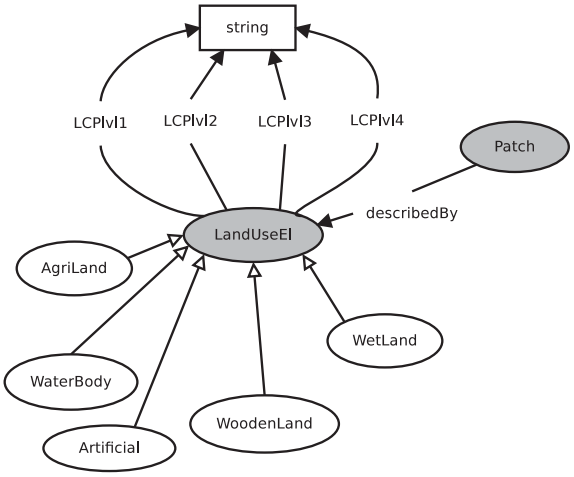

Figure 3: Portion of the LandUseElement hierarchy.

or maintenance (*Maintenance*) that is part of an intervention.

- *Patch* represents a small area of the map characterized by a specific land use; see below.

It is worth noting that a *Patch P* typically *belongsTo* an *EcologicalNetworkElement E*, and *E* is *madeOf* patches including *P*.

Figure 3 shows the *LandUseElement* hierarchy. Each class of the hierarchy is a singleton, containing exactly one representative object characterized by:

- A specific type of *land use* defined in the Land Cover Piemonte (LCP) cartography (Provincia di Torino, 2014): e.g., wetland (*WetLand*), wooden land (*WoodenLand*), and similar. The LCP defines 97 different types of land use at the most detailed level (*LCPlvl4*) of a hierarchy which includes 45 classes at level 3 (*LCPlvl3*), 15 classes at level 2 (*LCPlvl2*), and 5 general classes at the most general level *LCPlvl1*.
- Five evaluation criteria (OWL properties of *LandUseElement*, not shown in Figure 3) (Voghera and La Riccia, 2019):
  - *naturalness* (how close is the element to a natural environment);
  - *relevance* (how relevant it is for the conservation of the habitat);
  - *fragility* (how fragile the element is with respect to the anthropogenic pressure);
  - *extroversion* (how much pressure it can exert on the neighboring patches);
  - *irreversibility* (how difficult it would be to change the use of the element).

  The value for each criterion ranges from 1 to 5, and 1 represents the maximum value.

Each *Patch* of land is *describedBy* a LandUseElement; i.e., it is associated with a specific land use. For our current purposes, the description of land patches with the attributes we have just described is sufficient. In future work we may consider associating more refined information with patches, exploiting existing ontologies for modeling additional environment and ecologic concepts, e.g., ENVO (Buttigieg et al., 2016) and EcoCore (Buttigieg, 2018).

The *Intervention* and *Operation* hierarchies are intended for supporting the planning of improvements and expansions of the EN. At the current stage of our work we have modeled them with a few, simple concepts that are sufficient for our purposes. As we extend our work towards a full-fledged project planning support system (see section 7), we shall consider existing works such as (Lazoglou and Angelides, 2016; SDS Consortium, 2017), which aim at a detailed modeling of the concepts involved in spatial planning and spatial Decision Support Systems.

Compared to the ontology defined in (Torta et al., 2017; Torta et al., 2018), the described portion of the EN Ontology refines the class *ENElement* in the types of EN elements defined in (Bennett and Mulongoy, 2006). Moreover, it establishes a specific relation between EN Elements (e.g., *Corridor*) and the *Patches* they are made of. Furthermore, it specifies the land use characteristics of each patch by associating it with a *LandUseElement*.

3.2 Constraints Representation

Constraints in a geographical domain such as ENs must be able to express restrictions about the objects of the world modeled by the domain ontology, mixing logic, geometric, and numeric requirements; e.g., see (Louwsma et al., 2006). For example, they can refer to the allowed values of categorical attributes of areas, to the sum of the sizes of a set of areas, to the topological relations between pairs of areas, and so forth.

We have modeled constraints with a Constraint Ontology, whose classes refer to classes and properties of the EN Ontology[5]. We have defined several kinds of constraints, inspired by the kinds of constraints that may appear in a configuration knowledge base (Soininen et al., 1998; Stumptner et al., 1998; Felfernig et al., 2002). Figure 4 shows a portion of our Constraints Ontology, which is structured as a hierarchy rooted by class Constraint (in dark gray):

[5]In OWL, referring to classes and properties as values of other properties is a delicate point; see (W3C, 2017a). We currently avoid these difficulties by exploiting such references only in SPARQL queries.

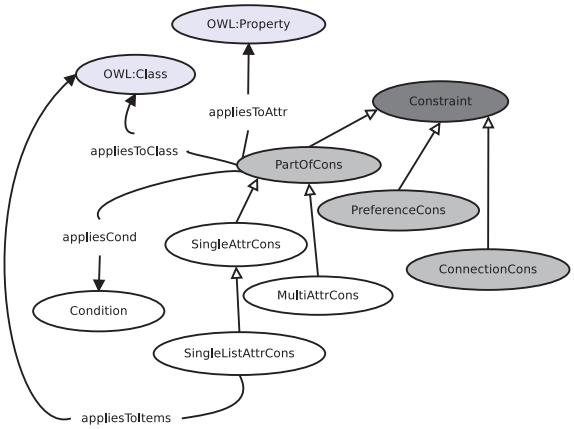

Figure 4: Diagrammatic representation of a portion of the constraints hierarchy.

- *PartOfConstraint* (PartOfCons) refers to a constraint that applies to one or more parts of a given object. Note that a part may be shared by more than one object; i.e., its semantics is similar to the *aggregation* of UML (Object Management Group, 2008).
  - *SingleAttributeConstraint* (SingleAttrCons) involves a single (part-of) attribute of the object.
  - *MultiAttributeConstraint* (MultiAttrCons) involves more than one (part-of) attribute of the object.
- *ConnectionConstraint* (ConnectionCons) refers to a constraint that applies to a relationship between more than one object.
- *PreferenceConstraint* (PreferenceCons) represents soft constraints, that augment regular constraints with functions to be optimized.

Let us focus on the *SingleAttributeConstraint*s (SACs). First of all, they refer to a single class (*appliesToClass*) and attribute of such class (*appliesToAttribute*). Note that an *attribute* of class *C* is an OWL property with class *C* as a possible domain. A special kind of SACs applies to attributes that are ordered lists of objects/values. In that case, a subset of the list elements to which the SAC applies can be identified, through the *appliesToItems* property. Currently, this property only allows the specification of the particular class of the items that should be affected by the constraint. In a future version, we may allow filtering conditions that select such elements based on their characteristics.

A SAC specifies a condition through the *appliesCondition* property; see Figure 5. The main distinction is between *AggregateCondition*s (AggregateCond, in gray), that specify restrictions on some aggregate quantity computed from the elements of a list attribute, and *IndividualCondition*s (IndividualCond) that apply to each value of a list attribute, or to the unique value of a scalar attribute. Individual conditions are built by composing *AtomicCondition*s (AtomicCond) into *CompositeCondition*s (CompositeCond) with the usual logic connectives *and*, *or*, *not*.

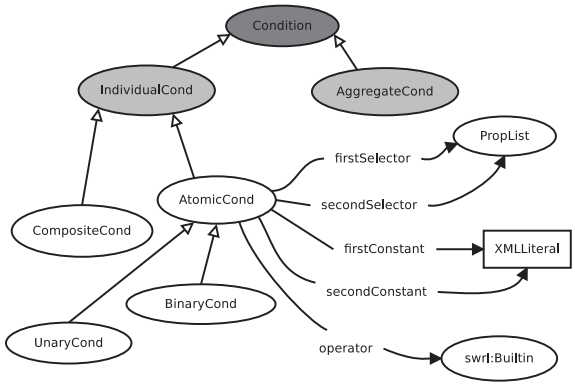

Figure 5: A portion of the condition hierarchy.

A *BinaryCondition* has two operands and an operator (*UnaryCondition*s have an empty *secondOperand*). The operands are either constant values (i.e., *XMLLiteral*s) or *selectors*. A selector is modeled as a list of properties: the list specifies a path from the element *E* that is being checked to a value that is reachable from *E* through the chain of properties in the list.

Note that the the main goal of the representation is the specification of various meta data about constraints; e.g., see the distinction between part-of, connection, and preference constraints. The ontology can be extended and refined as needed for expressing additional meta data. This is a key element for the development of reasoners that must automatically retrieve the suitable constraints to perform constraint solving, given the characteristics of the input problem.

Example 1. *Let us consider the* LandscapeCorridor *class in the EN ontology of Figure 2 (see also Figure 1). A specification taken from the guidelines for the Local EN implementation devised in project (Città Metropolitana di Torino, 2014) states that:*

> Corridors avoid areas with maximum irreversibility and areas with maximum extroversion.

A landscape corridor is therefore made of patches that must exhibit the specified characteristics. We can associate a suitable constraint with class LandscapeCorridor*. The constraint has the following traits:*

- *it is a* SingleListAttributeConstraint*, because the* patchList *property has a list value;*
- *it specifies an* IndividualCondition *that consists of the conjunction of two* BinaryCondition*s;*

- *the first* BinaryCondition *requires a non-maximum irreversibility; it specifies:*
  - *the* firstOperand *as a* PropertyList *containing the properties* describedBy *(from* Patch *to* LandUseElement*) and* irreversibility *(from* LandUseElement *to the value of the irreversibility criterion);*
  - *the* operator *as* swrlb:greaterThan*;*
  - *the* secondOperand *as the constant value 1.*
- *the second* BinaryCondition *is similar to the first one, but it requires non-maximum extroversion.*

3.3 Representation of Individual Information Items

The instances of the domain classes defined in the EN Ontology, such as the *Patch*es of land that form the map of a geographic area of interest, are stored in RDF format (W3C, 2017b) in a triple store that represents the knowledge base used by the system. The translation from input data-sets (typically available as ESRI shapefiles) to RDF triples is carried out by our data import functions; see Section 5.

Analogously, the triple store contains the instances of the constraints classes representing actual constraints that apply in the domain, such as the sample constraint described in Example 1.

4 GEOGRAPHIC REASONING

4.1 The Adjacency Graph Model

We aim to deal with geographic domains, as the Ecological Networks that motivated our work. Thus, it has been natural to base our reasoning system on a data model structured as a graph $\mathcal{G} = (N, E)$ where:

- the nodes N correspond to areas of a map;
- the arcs E connect nodes whose associated areas are adjacent in the map.

Figure 6 shows a map and the corresponding adjaciency graph. In the figure, the graph nodes corresponding to the areas of the map. For example, node *1* corresponds to an area that is adjacent to the areas corresponding to nodes *2*, *3*, *5*, and *6*. Note that, moving from the map to the graph model, some noise has been abstracted away; e.g., nodes *5* and *8* are connected although their areas do not actually touch each other. This abstraction is needed to deal with real-world, imperfect GIS data, and is a basic task that can be performed by our system. In particular, when a

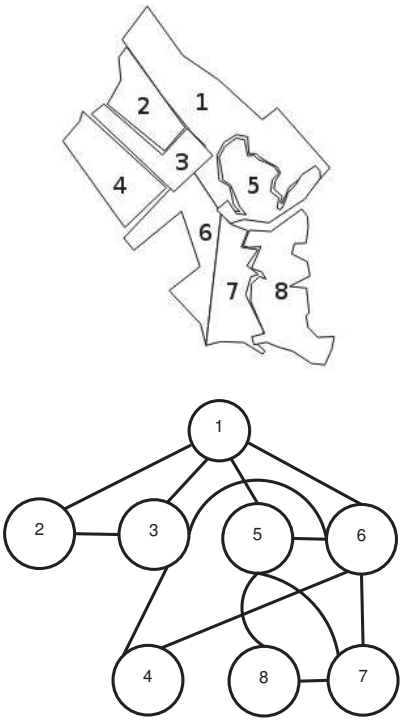

Figure 6: A map and its corresponding adjacency graph.

new area A is inserted into the knowledge base of the system, standard geometric algorithms can be used to:

(i) compute an expansion A' of A consisting of A with a border of a given thickness, and

(ii) determine the adjacent areas as the ones that intersect with such expanded area A'.

It is worth pointing out that, while GeoSPARQL provides a set of functions to compute the *Simple Features* topological relations it defines, those functions require that the involved geometries exactly satisfy the relation, e.g., an area *touches* another area iff they share some common points on their borders, but no internal points. This function is clearly too restrictive to determine the adjacency of two areas in a meaningful way for our purposes.

As a matter of fact, the graph model is able to easily accomodate more information than the association between nodes and areas:

- Each node $n \in N$ can have attributes representing meta-information about the area $\mathcal{A}(n)$ associated with n. In the EN domain, this may include:
  - the values of the evaluation criteria and the LCP levels of the *LandUseElement* describing $\mathcal{A}(n)$;
  - information such as the area size and perimeter of $\mathcal{A}(n)$;
  - the identity of the EN Element (e.g., *CoreArea*) to which $\mathcal{A}(n)$ belongs.

- Each arc $e = (n_i, n_j) \in E$ can have attributes that represent meta-information about the relationship between $\mathcal{A}(n_i)$ and $\mathcal{A}(n_j)$. For example:
  - the length of the perimeter shared by the two areas;
  - a numeric "cost" of going from $\mathcal{A}(n_i)$, to $\mathcal{A}(n_j)$ determined by the presence of a road/railway between the two areas.

4.2 Reasoning Tasks

Our model supports reasoning tasks based on three kinds of inputs:

- the OWL ontologies defining the domain concepts, the constraints hierarchy, and their relationships;
- the RDF data representing the instances of domain classes of the EN Ontology, and the instances of the constraints that apply to the specific domain;
- further requirements provided by the user to specify the desired reasoning task and its parameters.

Ideally, we would like that the system automatically extracts all the information needed to perform a requested task from the EN and Constraint Ontologies, and from the constraint instances, and use it to drive a generic reasoning engine to compute the answer from the RDF domain instances. However, such a generality would be extremely hard, if not impossible to achieve in practice. Instead, we equip the system with a pre-defined (but extensible) set of reasoning capabilities that can be reused for several specific tasks, and fill the details of specific reasoning task requests by exploiting all the input sources mentioned above.

Definition 1. *A* Reasoner $\mathcal{R}(\Omega, \Delta, \rho)$ *is a function that takes as inputs an OWL ontology* Ω*, a RDF graph* Δ *and a request* ρ *and:*

1. *extracts from* Δ *(driven by ontology* Ω*), a relevant set of constraints* $C = C_B \cup C_R$*;*
2. *builds an adjacency graph model* $\mathcal{G}$ *using the constraints in* C_B*;*
3. *performs a reasoning task on* $\mathcal{G}$ *using the constraints in* C_R*;*
4. *returns a result* α *that* answers *the request* ρ*, given* Ω *and* Δ*.*

The extraction of constraints is done by issuing SPARQL queries (W3C, 2017) on the RDF data Δ built from the vocabulary of ontology Ω; see the next section for an example. The retrieved constraints are then returned as data structures that can be directly interpreted by the reasoner in order to perform steps (2) and (3) above.

Specific tasks are requested by executing *Commands* that are translated to one or several invocations of the reasoners with specific values of ρ.

In the next subsection, we shall briefly describe two reasoners that we have implemented in order to support the automated identification, starting from geographic data about an area, of core areas, and corridors, according to the specifications given in the EN Ontology.

4.3 Identifying Core Areas and Corridors

Reasoner #1. The $\mathcal{R}_{PATH}$ reasoner is meant to receive through the request θ two identifiers id_s and id_e of *Patch*es or elements that aggregate *Patch*es (i.e., EN elements), and the name of a property *prop* that is a list of *Patch*es or elements that aggregate *Patch*es. It then computes a path of adjacent elements from element id_s to id_e taking into account the constraints associated with property *prop*.

This reasoner can be used to implement the command *BUILD(LandscapeCorridor, id_s, id_e)* which assigns id_s, id_e to the *from* and *to* attributes of *LandscapeCorridor*, and computes the value of the *patchList* attribute by invoking $\mathcal{R}_{PATH}$ with $\theta = (id_s, id_e, patchList)$.

1. First of all, the reasoner issues a number of SPARQL queries to retrieve the *SingleListAttributeConstraint C* associated with *patchList* (described in Example 1).
2. Then, it builds an adjacency graph $\mathcal{G}$ in such a way that the nodes of $\mathcal{G}$ are associated to patches that satisfy *C*; i.e., they have non-maximum irreversibility and extroversion.
3. Finally, the reasoner applies a simple path-finding algorithm based on the well-known A^* algorithm (Dechter and Pearl, 1985) to identify a corridor between the id_s and id_e elements, if any.

Reasoner #2. The $\mathcal{R}_{CLUST}$ reasoner is meant to receive, through request θ, the identifier *id* of a *Patch*, and the name of a property *prop* that is a one-to-many relationship from a class and *Patch*. Given these inputs, the reasoner computes a clustering of patches that satisfy *prop*, starting from *id*, by taking into account the constraints associated with property *prop*.

The reasoner can be used to implement the command *BUILD(CoreArea, id)*, which computes the value of the *madeOf* attribute by invoking $\mathcal{R}_{CLUST}$ with $\theta = (id, madeOf)$. Note that the *madeOf* attribute of *CoreArea* is associated with a constraint that specifies the characteristics of the patches *P* that can be included in a core area. These are the patches with high or

medium *ecological functionality*, which is a function of the *naturality* and *relevance* of the *LandUseElement* that describes patch *P*.

5 IMPLEMENTATION

We have implemented the model described in the previous sections as a Java library consisting of the following modules:

- *data-import* contains functions supporting the import/export of shape files to/from a triple store (e.g., Parliament (Battle and Kolas, 2012)), the pre-processing and optimization of the geometries associated with geo-SPARQL *Feature*s, and the transfer of RDF triples between disk and the in-memory model of the Jena library (https://jena.apache.org/) used to query the triple store;
- *reasoning* contains the functions for the creation of the adjacency graph data model. Moreover, it collects all the specific reasoners provided by the system (currently, the $\mathcal{R}_{PATH}$ and $\mathcal{R}_{CLUST}$ reasoners described above);
- *commands* implements the parsing of commands (currently, the two forms of the *BUILD* command described above) and interfaces with the *reasoning* and *data-import* functions to execute them;
- *shared* provides the definitions and implementations of elements relevant across the other modules, such as the geometric feature and triple store manager, as well as utility functions used by the other modules.

By exploiting the *data-import* module, we have populated the Parliament triple store with 395 patches defined in the shape files of a portion of map situated at the north of the Italian city of Turin. We have then used the implementation of the $\mathcal{R}_{CLUST}$ reasoner contained in the *reasoning* module to generate the Core Areas as clusters of patches with given characteristics. The reasoner has generated 74 clusters. Finally, we have used the implementation of the $\mathcal{R}_{PATH}$ reasoner to generate a number of landscape Corridors between pairs of Core Areas specified by us.

6 RELATED WORK

To the best of our knowledge, the only work pursuing the automation of a task needed for EN project planning is the one described in (Torta et al., 2017; Torta et al., 2018). In that work, the authors defined a machine readable specification of ENs to support EN validation. They introduced an ad-hoc constraint satisfaction language for the verification of EN constraints on a geographical area; e.g., to check whether a certain area, identified as a Buffer zone in a predefined EN, complies with the definition given in the EN specifications, or not. While compliance verification with respect to an *pre-defined* EN can be useful to help the human planners in an EN planning task, our aim is to take a step further, by addressing the automatic suggestion of suitable aggregations of land patches into elements of the EN. Moreover, we aim at automatically suggesting *modifications* to existing elements needed by the urban design projects, through suitable interventions. Even though the current implementation of our model supports a limited kind of solutions to tasks only involving two types of elements of ENs (e.g., the suggestion of modifications is out of the scope of the present paper), the proposed approach is designed to support full fledged implementations of creation and modification tasks.

Another key aspect of our method is the representation of constraint types as classes of an OWL ontology (the Constraints Ontology). In this way, we employ a single, well-known standard for the representation of semantic knowledge, and we avoid the introduction of a new language, requiring ad-hoc tools for managing constraint information. Even more importantly, an ontology of constraints supports a detailed *description* of the different kinds of constraints (e.g., soft and hard, part-of and relational, aggregation and individual) by qualifying their scope, purpose, relationships, and so forth. This meta-information about constraints enables the reasoners to retrieve the constraints that are relevant for the task at hand, and use these constraints in the correct places and in the correct way. This is the reason why we did not adopt any of the existing ways to represent constraints in the Semantic Web, including *rule languages* such as SWRL (Semantic Web Rule Language) (Horrocks et al., 2004) and RuleML (Boley et al., 2001), or any logical/functional languages such as CIF (Constraint Interchange Format) (Gray et al., 2001). Those languages allow the definition of constraints as generic rules or logic formulas, without characterizing their properties as needed to automatically retrieve constraints and apply them for specific reasoning tasks.

Regarding the current implementation of the $\mathcal{R}_{PATH}$ reasoner, it is worth noting that path finding has been explored in recommender systems research to suggest travel paths suiting individual preferences; e.g., the shortest path between two endpoints, or a path maximizing pleasure, calm, or other numerical properties of an area (Quercia et al., 2014). Given

a graph representing the travel map of a geographical area, those works choose the path(s) composed of road segments which, globally, maximize one or more measures associated to the selected properties. These approaches solve a specific task taking into account a pre-defined set of constraints. Differently, we aim at supporting multiple reasoning tasks and retrieving relevant constraints from a semantic knowledge base (e.g., an RDF store) by exploiting their description as classes of an OWL ontology.

7 FUTURE WORK

The work presented in this paper builds a solid ground, both from the modeling and reasoning points of views, for supporting a user in the project planning for a geographical domain such as the one of Ecological Networks. However, there are several directions in which our work may be extended to provide a full fledged decision support system that human users can effectively use in the definition of public policies for land use. For instance, we plan to:

- extend the Constraint Ontology and refine it to specify more types of constraints, especially *soft* constraints for guiding the reasoners to compute solutions that maximize some preference criteria; and sophisticated *geometric* constraints about the shapes, sizes, etc. of given areas;
- extend our reasoning framework with the ability of exploiting soft constraints, and with additional reasoners (e.g., for proposing maintenance and modification interventions on an EN with the goal of expanding or improving it);
- create a GUI-based application to support the mixed-initiative interaction between the human user and the reasoning system to jointly perform project planning tasks.

8 CONCLUSIONS

We described a semantic framework supporting various types of constraint-based reasoning tasks on a geographic domain, from the basic validation of conditions in a geographic area, to the suggestions for defining and modifying EN elements. Different from other constraint-based geographic reasoners, our model represents both the geographic domain and its constraints in OWL ontologies. This approach has several advantages: first of all, it does not introduce any ad-hoc language for the management of constraints. Second, it fully exploits the knowledge representation and reasoning interoperability provided by semantic languages for knowledge sharing and for data/application interoperability. Third, it opens the avenue to the classification of constraints for their automated management within reasoners able to adapt to solve a range of reasoning problems.

As a test-bed for our framework, we considered the domain of Ecological Networks. Whereas, at the current stage, we implemented reasoning about ENs as a stand-alone model, the main motivation and application of our work lies in its possible integration within Participatory Geographical Information Systems (PGIS, (Sun and Li, 2016)), in order to support online interaction with stakeholders in inclusive processes aimed at collecting feedback and project proposals from stakeholders.

ACKNOWLEDGEMENTS

This work is partially funded by project MIMOSA (MultIModal Ontology-driven query system for the heterogeneous data of a SmArtcity, “Progetto di Ateneo Torino_call2014_L2_157”, 2015-17), and by “Ricerca Locale” of the University of Torino. We thank Adriano Savoca for his precious contributions to the initial steps of this work.

REFERENCES

Baglatzi, A., Kokla, M., and Kavouras, M. (2012). Semantifying openstreetmap. In *Proc. of Terra Cognita Workshop on Foundations, Technologies and Applications of the Geospatial Web - CEUR Workshop Proceedings, Vol. 901*, page paper 4. CEUR.

Battle, R. and Kolas, D. (2012). Enabling the geospatial Semantic Web with Parliament and GeoSPARQL. *Semantic Web*, 3(4):355–370.

Benedict, M. and McMahon, E. (2002). *Green Infrastructure: smart conservation for the 21st century*. Watch Clearinghouse Monograph Series. Sprawl.

Bennett, G. and Mulongoy, K. (2006). Review of experience with ecological networks, corridors and buffer zones. *Technical Series*, 23.

Bennett, G. and Wit, P. (2001). The development and application of ecological networks: a review of proposals, plans and programmes. *AIDEnvironment*.

Boley, H., Tabet, S., and Wagner, G. (2001). Design rationale of ruleml: A markup language for semantic web rules. In *Proceedings of the First Int. Conf. on Semantic Web Working*, pages 381–401. CEUR-WS. org.

Buttigieg, P. L. (2018). Ecology Core Ontology. https://github.com/EcologicalSemantics/ecocore. Accessed: 2018-07-28.

Buttigieg, P. L., Pafilis, E., Lewis, S. E., Schildhauer, M. P., Walls, R. L., and Mungall, C. J. (2016). The environment ontology in 2016: bridging domains with increased scope, semantic density, and interoperation. *Journal of biomedical semantics*, 7(1):57.

Città Metropolitana di Torino (2014). Misura 323 del PSR 2007-2013 (in Italian). http://www.cittametropolitana.torino.it/cms/territorio-urbanistica/misura-323/misura-323-sperimentale.

Codescu, M., Horsinka, G., Kutz, O., Mossakowski, T., and Rau, R. (2011). OSMonto – an ontology of openstreetmap tags. In *Proc. State of the map Europe (SOTM-EU) 2011*, pages 55–64, Vienna, Austria.

Council of Europe (2000). General guidelines for the development of the pan-european ecological network. *Nature and environment*, 107.

Dechter, R. and Pearl, J. (1985). Generalized best-first search strategies and the optimality of a*. *Journal of the ACM*, 32(3):505536.

Fath, B., Sharler, U., Ulanowicz, R., and Hannon, B. (2007). Ecological network analysis: network construction. *Trends in Ecology & Evolution*, 208:49–55.

Felfernig, A., Friedrich, G., Jannach, D., and Zanker, M. (2002). Configuration knowledge representation using uml/ocl. In *International Conference on the Unified Modeling Language*, pages 49–62. Springer.

Fonseca, F., Egenhofer, M., Agouris, P., and Câmara, G. (2002a). Using ontologies for geographic information systems. *Transactions in GIS*, 3:231–257.

Fonseca, F., Egenhofer, M., C.A. Davis, and Câmara, G. (2002b). Semantic granularity in ontology-driven geographic information systems. *Annals of Mathematics and Artificial Intelligence*, 36(1-2):121–151.

GeoNames.org (2018). Geonames mappings ontology. http://www.geonames.org/ontology/mappings_v3.01.rdf.

Gobluski, A., Westlund, E., Vandermeer, J., and Pascual, M. (2016). Ecological networks over the edge: Hypergraph trait-mediated indirect interaction (TMII) structure. *Trends in Ecology & Evolution*, 31(5):344354.

Gray, P., Hui, K., and Preece, A. (2001). An expressive constraint language for semantic web applications. In *E-Business and the Intelligent Web: Papers from the IJCAI-01 Workshop*, pages 46–53.

Hall, M. and Mandl, P. (2006). Spatially extended ontologies for a semantic model of harmonised landuse and landcover information.

Horrocks, I., Patel-Schneider, P., Boley, H., Tabet, S., Grosof, B., and Dean, M. (2004). SWRL: A semantic web rule language combining OWL and RuleML. *W3C Member submission*, 21.

Janowicz, K., Scheider, S., Pehle, T., and Ha, G. (2012a). Geospatial semantics and linked spatiotemporal data past, present, and future. *Semantic Web - On linked spatiotemporal data and geo-ontologies*, 3(4):321–332.

Janowicz, K., Scheider, S., Pehle, T., and Ha, G. (2012b). Linkedgeodata: A core for a web of spatial open data. *Semantic Web Interoperability, Usability, Applicability*, 3(4):333–354.

Jongman, R. (1995). Nature conservation planning in europe: developing ecological networks. *Landscape and urban planning*, 32(3):169–183.

Lazoglou, M. and Angelides, D. C. (2016). Development of an ontology for modeling spatial planning systems. *Current Urban Studies*, 4(02):241.

Louwsma, J., Zlatanova, S., van Lammeren, R., and van Oosterom, P. (2006). Specifying and implementing constraints in GIS - with examples from a geo-virtual reality system. *GeoInformatica*, 10(4):531–550.

Lurgi, M. and Robertson, D. (2011). Automated experimentation in ecological networks. *Automated Experimentation*, 3(1).

Montenegro, N., Gomes, J., Urbano, P., and Duarte, J. (2011). An owl2 land use ontology: Lbcs. In *International Conference on Computational Science and Its Applications*, pages 185–198. Springer.

Object Management Group, O. (2008). Unified Modeling Language (UML. http://www.uml.org.

OGC (2012). Geosparql vocabulary. http://schemas.opengis.net/geosparql/1.0/geosparql_vocab_all.rdf.

Open Geospatial Consortium et al. (2011). *OpenGIS Implementation Standard for Geographic information-Simple feature access-Part 1: Common architecture*.

OpenStreetMap Contributors (2017). Openstreetmap. https://www.openstreetmap.org.

Palacio, D., Derungs, C., and Purves, R. (2015). Development and evaluation of a geographic information retrieval system using fine grained toponyms. *Journal of Spatial Information Science JoSIS*, 11:1–29.

Pilosof, S., Porter, M., Pascual, M., and Kefi, S. (2017). The mulutilayer nature of ecological networks. *Nature ecology&evolution*, 1:article 101.

Provincia di Torino (2014). Linee guida per le reti ecologiche (in italian). http://www.provincia.torino.gov.it/territorio/file-storage/download/pdf/pian_territoriale/rete_ecologica/lgsv_lgre.pdf.

Quercia, D., Schifanella, R., and Aiello, L. (2014). The shortest path to happiness: Recommending beautiful, quiet, and happy routes in the city. In *Proc. of the 25th ACM Conference on Hypertext and Social Media*, HT '14, pages 116–125, New York, NY, USA. ACM.

SDS Consortium (2017). Spatial decision support ontology. http://sdsportal.sdsconsortium.org/ontology/.

Soininen, T., Tiihonen, J., Männistö, T., and Sulonen, R. (1998). Towards a general ontology of configuration. *Ai Edam*, 12(4):357–372.

Stumptner, M., Friedrich, G. E., and Haselböck, A. (1998). Generative constraint-based configuration of large technical systems. *AI EDAM*, 12(4):307–320.

Sun, Y. and Li, S. (2016). Real-time collaborative GIS: a technological review. *ISPRS Journal of Photogrammetry and remote sensing*, 115:143–152.

Torta, G., Ardissono, L., Corona, M., La Riccia, L., Savoca, A., and Voghera, A. (2018). GeCoLan: a constraint language for reasoning about ecological networks in the semantic web. In *Knowledge Discovery, Knowledge Engineering and Knowledge Management, IC3K 2017, Revised Selected Papers*, page

to appear. Springer-Verlag, Berlin Heildelberg New York.

Torta, G., Ardissono, L., Savoca, A., Voghera, A., and Riccia, L. L. (2017). Representing ecological network specifications with semantic web techniques. In *Proc. of 9th Int. Joint Conf. on Knowledge Discovery, Knowledge Engineering and Knowledge Management (KEOD 2017)*, pages 86–97, Funchal, Madeira, Portugal. SCITEPRESS.

Ulanowicz, R. (2004). Quantitative methods for ecological network analysis. *Computational biology and chemistry*, 28:321339.

Van Hage, W. R., Malaisé, V., Segers, R., Hollink, L., and Schreiber, G. (2011). Design and use of the simple event model (sem). *Web Semantics: Science, Services and Agents on the World Wide Web*, 9(2):128–136.

Voghera, A. and La Riccia, L. (2019). Ecological networks in urban planning: Between theoretical approaches and operational measures. In Calabrò, F., Della Spina, L., and Bevilacqua, C., editors, *New Metropolitan Perspectives*, pages 672–680, Cham. Springer International Publishing.

W3C (2017a). Representing classes as property values on the semantic web. https://www.w3.org/TR/2005/NOTE-swbp-classes-as-values-20050405/.

W3C (2017b). Resource description framework (RDF). https:// www.w3.org/RDF/.

W3C (2017). SPARQL query language for RDF. https:// www.w3.org/TR/rdf-sparql-query/.

W3C (2017). Web ontology language (OWL). https:// www.w3.org/OWL/.