

Spatial Planning on the Semantic Web

Radboud Winkels¹, Rinke Hoekstra^{1,2}, and Erik Hupkes¹

¹ Leibniz Center for Law, Universiteit van Amsterdam
Kloveniersburgwal 48, 1012 CX, Amsterdam, The Netherlands
{winkels,hoekstra,ehupkes}@uva.nl

² AI Department, Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081HV, Amsterdam, The Netherlands
hoekstra@few.vu.nl

Abstract. Land use regulations are an important but often underrated legal domain. In densely populated regions such as the Netherlands, spatial plans have a profound impact on both (local) governments and citizens. This paper describes our work on a ‘Legal Atlas’. Using Semantic Web technology we combine distributed geospatial data, textual data and controlled vocabularies to support users in answering questions such as “What activity is allowed here?”. Spatial norms are represented using OWL 2 in a way that enables intuitive visualisation of their effects: map based legal case assessment. Users can represent a (simple) case by selecting or drawing an area on the map. Given a designation for that area, the system can assess whether this is allowed or not. The same solution also enables the comparison of two or more sets of spatial norms that govern the same region.

1 Introduction

Land use regulations are an important but often underrated legal domain. Especially in densely populated regions such as the Netherlands, spatial plans have a profound impact on both (local) governments and citizens alike. They determine whether building permits are granted, businesses can expand, roads can be widened, housing projects can be built, and even determine the chance of survival of endangered species. Although spatial regulations can be characterised as normal regulations – given their applicability within a spatially restricted jurisdiction (e.g. a country) – spatial plans regulate at a much more fine-grained level, where individual norms may apply only to specific locations or regions (often called zones). In practice this has led to a situation where the annotated map of a spatial plan is the dominant source of law, rather than the textual description of the plan.

Over the past decade, internet technology significantly improved the accessibility of legal sources. This development not only meant the construction of internet portals for accessing legal texts, but also involved a standardisation of the structural description of legal sources in terms of both web standards – such as XML – and interchange standards – such as CEN MetaLex [1, 2].¹ Given the

¹ See <http://www.metalex.eu>

prominence of maps in spatial plans, in order to improve access to these regulations a combination of existing technology for disclosing legal texts with that currently available in geographical information systems seems inevitable. Such a combination allows the connecting of textual descriptions as in traditional legal sources, to the more object oriented representation of the world employed by the maps of zoning plans. The Legal Atlas [3, 4] was the first system that combined MetaLex encoded texts with the corresponding maps expressed in GML² by using Semantic Web technology, RDF in particular.³ In this paper we show that this proves a solid approach for extending the functionality of Legal Atlas well beyond straightforward information integration; i.e. to support not just concept-based information retrieval, but rather map-based *normative reasoning*.

Legal Atlas can serve several use cases and scenarios:

Assessment – A citizen or company or a civil servant consults the regulations to see whether a particular plan at a specific location is allowed. In case there is the option to compensate somewhere else in order for the original plan to be approved, the system should indicate this on the map. In the spatial planning domain it is sometimes allowed to compensate for violating a norm in one area in another area, e.g. compensating the building of houses in a park by turning another area into a new park

Planning – A citizen or company consults the regulations to see *where* a particular plan is allowed.

Evaluation – Someone wants to or needs to check the consistency of two sets of spatial norms. The State Council in the Netherlands for instance, may need to check whether a zoning plan of a city council fits the overall zoning plan of the province. A second type of evaluation by a legislative drafter for instance may be to see what the consequences of a new set of norms might be in practice. Suppose the idea is to forbid “Dutch coffee shops”⁴ in the vicinity of schools (an actual case in Amsterdam at the moment). If we have access to a database of all schools in a certain area with geographical location, the consequence of various implementations of these norms (e.g. not within a radius of 500 meters, or 1000 meters etc.) can be plotted on a map. We can for instance see if any areas remain where coffee shops would be allowed at all (cf. the “Planning” scenario above). If we also have a database of all existing coffee shops in the area, we can determine the consequences for these establishments.

2 Requirements

It is clear that Legal Atlas has to deal with the intrinsic heterogeneity of the information it discloses: spatial plans are not just documents, but they are closely tied to maps. These maps are created using advanced Geographic Information Systems (GIS) technology, and are ideal candidates for standardisation. Geographical standardisation efforts focus primarily on exchange formats, such as

² GML: the open Geography Markup Language, see <http://www.opengis.net/gml/>.

³ RDF: the Resource Description Framework, see <http://www.w3.org/RDF>

⁴ Well known for other products than coffee.

GML, KML and ESRI Shapefiles,⁵ but also on standard web-based service protocols, e.g. the Web Map Service (WMS) and Web Feature Service (WFS) protocols.⁶ To comply with these standards, geographical information servers are required to implement basic facilities for spatial reasoning, such as determining the overlap between two polygons (regions).

The free availability of web-based application programming interfaces, most notably the Google Maps API, has exposed this technology to the open source community at large, and has already led to various mash-ups of incredibly diverse data and maps: geo-tagging of photos allows us to search for pictures of food in Thailand, we can view map overlays of traffic conditions all over the world, and even consult weather radar imagery. In all of these cases, the key means to tie together the two types of information – photos and maps, traffic data and maps, or legal texts and maps – is to determine the overlap of data that is meaningful in both domains. Where for traffic data, the location (stretch of road) and traffic density are obvious candidates, in spatial planning one typically depicts the category of land use applicable to each area. Therefore, to connect maps to texts, and vice versa, they should be annotated using the same set of metadata [3].

However, for any given region a multitude of different regulations applies. These regulations may be issued by different government bodies with overlapping jurisdiction; government bodies that each maintain a categorisation specifically targeted towards the types of maps it produces. This is problematic for a number of reasons. Firstly, users of a system that discloses these regulations may not care for, or even be aware of the borders and interaction between these jurisdictions. The system should therefore be able to present all applicable regulations in a *transparent* fashion, on a single map. Secondly, if for instance land use in regulations issued by *adjoining* municipalities is categorised differently, the land use ‘regime’ in these municipalities cannot be *compared*.

Comparability of regulations is not just beneficial to individuals or businesses that need to decide where to conduct a certain activity, it is a crucial ingredient for the harmonisation of these regulations. Ideally, categorisations are therefore shared among governments, even when these governments issue their respective regulations in different languages, as is e.g. the case within the EU.

The need for shared categorisation of types of land use has been acknowledged in a number of standardisation initiatives. For instance, the IMRO 2006 standard was issued by the Dutch ministry of spatial planning and environment, and is currently obligatory for all urban planning by municipalities in the Netherlands.⁷ IMRO specifies a strict categorisation system for all sanctioned types of land use

⁵ KML, the Keyhole Markup Language is Google’s variant of GML, see http://earth.google.com/intl/en/userguide/v4/ug_kml.html. ESRI is a leading developer of GIS software, see <http://www.esri.com/library/whitepapers/pdfs/shapefile.pdf> for the technical description of ESRI Shapefiles.

⁶ Both WMS and WFS are standards of the Open Geospatial Consortium, see <http://www.opengeospatial.org/>.

⁷ IMRO: Informatiemodel Ruimtelijke Ordening (information model for spatial planning), see <http://www.geonovum.nl/informatiemodellen/imro-2006/> (Dutch).

at multiple levels of government, and even prescribes the colours to be used on compliant maps. At the European level, the INSPIRE initiative similarly aims to standardise the spatial information published by governments across all member states.⁸ Although not specifically targeted to spatial information, the GEMET thesaurus provides a huge multilingual vocabulary for information related to the environment.⁹ The 5000+ terms defined in GEMET are organised as a hierarchy, and available in 27 languages, which makes it ideally suited for cross language mapping of terms.

Where IMRO and INSPIRE introduce a *top-down* standard for direct exchange of information between organisations within the *same* domain, GEMET functions rather as *umbrella* that provides a unifying framework for information exchange between organisations that deal with *different* (sub) domains. For instance, GEMET identifies 40 themes ranging from ‘forestry’ to ‘physics’, that are of relevance to all levels of government. An integrated, transparent unlocking of content tagged using such different schemes must take into account a *mapping* between them. Of course, an additional challenge is that different categorisation schemes may rely on different technology. Although GEMET is defined as a SKOS vocabulary, IMRO is specified as an XML Schema, and accompanying UML diagrams.¹⁰

A design goal of both IMRO and INSPIRE is that maps produced at different levels of government may be depicted as layers on a single map. This poses interesting opportunities from a legal perspective. Quite distinct from ordinary maps, governments issue spatial plans not to *describe* the existing situation on the ground, but rather to *prescribe* the restrictions and rights associated with geospatial objects. [4] identify two important characteristics of a map-based interface to regulations. Firstly, hierarchical relations between authorities are mirrored in *spatial inclusion* relationships. We can exploit knowledge of the regulatory bodies involved to resolve normative conflicts between overlapping regions by applying the *Lex Superior* priority ordering of norms, which states that ‘higher’ norms prevail over ‘lower’ norms (e.g. national legislation defeats local regulations if these conflict).

Secondly, geospatial *adjacency* information can be used to determine indirect effects; rules and restrictions can be naturally grouped by area of effect, e.g. to prevent the development of an industrial zone right next to a nature reserve. We can use adjacency information to apply a larger number of relevant norms to an area, than we would be able to without this information. This information may be a valuable contribution to cross-border decision making processes.¹¹

⁸ INSPIRE: Infrastructure for Spatial Information in the European Community, see <http://inspire.jrc.ec.europa.eu/>.

⁹ GEMET: GEneral Multilingual Environmental Thesaurus, see <http://www.eionet.europa.eu/gemet>.

¹⁰ SKOS: Simple Knowledge Organization System, see <http://www.w3.org/TR/skos-reference>.

¹¹ This is in fact the main focus of the FEED project.

Maps thus provide intuitive handles for evaluating the normative content of spatial plans. However, although even the most rudimentary geoservers (e.g. MySQL with spatial extensions¹²) provide some form of spatial reasoning, no such functionality exists for the normative aspects of spatial plans. The Legal Information Server (LIS) of [5] was the first attempt to provide normative reasoning as a *service*. Its main task was to *assess* whether some situation description is allowed or disallowed given a set of norms [6]. Where the LIS implementation depended on a custom-built representation in Prolog, the more recent implementation of [7] represent regulations using OWL 2 DL [8].¹³

Arguably, the use of a restricted language such as OWL 2 DL poses some problems for the representation of regulations. Firstly, certain aspects of legal reasoning may be hard or impossible to represent using a language that depends on monotonic reasoning. Especially considering the commonly held claim that legal reasoning is in principle *defeasible* [9]: reasoning follows a *dialectic process of argumentation* by which contrary positions are continuously attacked and revised. The use of a defeasible logic serves a practical purpose in that it enables us to circumvent situations where normative conflicts give rise to logical contradictions: considered separately, each defeasible state is consistent, and without taking their defeasibility into account, the whole would be inconsistent. Monotonic approaches consider each state separately, depending on knowledge-base updates for moving from one state to another. Given the limited task at hand – normative assessment – and a user interface where users may ‘play around’ with possible input, the restricted expressiveness of DL is instrumental to our purposes in that it offers guarantees to response times [10]. Furthermore, the work of [7] has shown that a significant portion of exceptions between norms, such as *lex specialis*, can be dealt with without resorting to defeasible representations.

A second potential problem for using OWL 2 DL concerns not the complexity of legal reasoning, but rather the complexity of the *world* that is governed by law [11]. To retain decidability, OWL 2 DL is restricted to models that have the ‘tree property’: situations that describe complex configurations of multiple objects can be only approximately defined in DL [12, 11]. In the case of spatial information, this may pose problems in the description of relations between areas. Consider for instance the definition of a body of water, e.g. a river R , that separates two areas of land, L_1 and L_2 . In OWL 2 DL, and other decidable description logics, R cannot be defined as it is impossible to exclude in its definition a situation where L_1 and L_2 overlap.

However, the domain itself is otherwise at a relatively high level of abstraction that excludes the complex structures found in domains such as biology and engineering. Not only can spatial reasoning be delegated to a geoserver; the land use categories defined in IMRO and INSPIRE provide a standard intermediary for communicating normative content (see Section 4).

¹² See <http://www.mysql.com>.

¹³ OWL 2 is an extension of the OWL Web Ontology Language, currently in Proposed Recommendation stage of the W3C standardisation procedure, see <http://www.w3.org/2007/OWL>.




Fig. 1. Relations between maps, metadata and texts in Legal Atlas.

A final requirement for the system described here, is that it should cater for not just heterogeneity in *content*, but also in *location*. We have argued that the regulations issued by different authorities are to be made visible on the map to improve transparency and comparability. However, these regulations are also *maintained* by these authorities, and are continuously subject to change. It is important therefore that the spatial and semantic *representation* of regulations should likewise be maintained by their respective owners. The system should therefore be able to cope with distributed content and semantics: a prime use case for Semantic Web technology [13].

3 Maps, Metadata and Texts in a Prototype System

We represent the relation between texts, maps, and metadata as depicted in Figure 1. Texts are annotated using `metalex:region` properties that point to a region. The `la:Region` class describes all individual regions in a spatial plan. Regions are related to polygons stored in a GeoServer instance via the `geo:id` property, and to an individual of type `la:Land_Use` via the `la:land_use` property. Contrary to the initial approach [3], these individuals are described as SKOS concepts, i.e. they are simultaneously of type `skos:Concept` and of type `la:Land_Use`. Consequently, land use categorisations such as captured by IMRO and GEMET are represented as separate `skos:Schema` hierarchies of `skos:Concepts`. E.g. a model that allows only the use of land use categories from GEMET and IMRO defines the class `la:Land_Use` as those individuals of type `skos:Concept` that have a `skos:inScheme` relation with either `gemet:gemet` or `imro:imro2006`. See also Figure 2.

This approach leaves freedom to specify restrictions on combinations of land use independently from either IMRO or GEMET, which benefits maintainability of our models. The next sections present the architecture of a prototype implementation, and describe how the representation of regions and land use has been incorporated in the specification of spatial norms.

The current version of Legal Atlas is a prototype system that allows users to query for potential conflicts.¹⁴ Given a set of land use categories, the system

¹⁴ See <http://feed.leibnizcenter.org>.

Correspondingly we define two regions, `ex:IndustryRegion` and `ex:NatureRegion`, both of which are subclasses of `la:Region`:

$$\begin{aligned} \text{ex:IndustryRegion} &\equiv (\text{la:land_use } \mathbf{some} \text{ ex:Industry}) \sqcap \text{la:Region} \\ \text{ex:NatureRegion} &\equiv (\text{la:land_use } \mathbf{some} \text{ ex:Nature}) \sqcap \text{la:Region} \end{aligned}$$

We map the corresponding land use categorisations from IMRO by assigning them as individuals belonging to the respective classes. This representation has a number of benefits, the most important being that our representation of land use is independent of the categorisation scheme adopted. For instance, we can map *gemet:industry* to `ex:Industry` in the same way, allowing us to specify norms on both IMRO and GEMET encoded maps.

`norm:NoOverlapIN` is a norm that states that an overlap between these two types of regions is not allowed. In simplification of the approach presented in [7], where norms are kept disjunct from the domain being governed, this norm is represented as a subclass of `ex:IndustryRegion` with the restriction that the region overlaps with a region of type `ex:NatureRegion`. This norm is simultaneously defined as a subclass of `norm:ConflictRegion`:

$$\begin{aligned} \text{norm:NoOverlapIN} &\equiv \text{ex:IndustryRegion} \sqcap \text{la:overlaps } \mathbf{some} \text{ ex:NatureRegion} \\ &\sqsubseteq \text{norm:ConflictRegion} \end{aligned}$$

As an illustration of the representation, consider the following spatial plan, specifying a single region with land use *imro:nature*:

$$\begin{aligned} \text{ex:region_nature} &\in \text{la:Region} \\ \text{la:land_use}(\text{ex:region_nature}, \text{imro:nature}) \end{aligned}$$

A user draws a region *ex:region_industry* and specifies its land use to be *imro:industry*. The system queries the GIS server to find all regions that *ex:region_industry* overlaps with. Suppose, in this case we find `ex:region_nature`, we then update the repository with the assertion:

$$\text{la:overlaps}(\text{ex:region_industry}, \text{ex:region_nature})$$

The third step is to perform OWL 2 DL realization on the knowledge base. This will infer that *ex:region_industry* \in `ex:IndustryRegion` and *ex:region_nature* \in `ex:NatureRegion` (the dotted lines in 3). However, *ex:region_industry* also meets the requirements of `norm:NoOverlapIN`, and we can therefore infer *ex:region_industry* \in `norm:NoOverlapIN`.

Finally, the system will gather all individuals of the class `norm:ConflictRegion` using a simple SPARQL query:¹⁸

```
SELECT ?region
WHERE { ?region rdf:type norm:ConflictRegion . }
```

Because `norm:NoOverlapIN` is a subclass of `norm:ConflictRegion`, *ex:region_industry* will be bound to the `?region` variable. The system will bring this fact to the attention of the user by highlighting his region on the map.

¹⁸ <http://www.w3.org/TR/rdf-sparql-query/>

5 Conclusions and Discussion

This paper presents an approach for specifying spatial norms using Semantic Web technology that enables an intuitive way of visualising their effects: map based legal case assessment. Users can see what is allowed and what not in specific areas on the map, they can represent a (simple) case by selecting or drawing an area on the map. Given a designation for that area, they can have the system assess whether this is allowed or not. The same solution also enables the comparison of two or more sets of spatial norms that govern the same region, e.g. coming from a municipality and the province it is part of. We have demonstrated a practical use of the case assessment method specified in [7] using OWL 2 DL, and presented a prototype system that provides a partial implementation of the approach.

The simple assessment of cases or the detection of conflicts does not constitute full legal reasoning, but Legal Atlas gives users the opportunity to experiment with their plans and prevailing norms. By using various queries they can find out where their plans might be allowed and where they will meet resistance. A next step will be explaining or justifying the results. Partly this can be done by referring to the original sources of law through the links to the MetaLex representation (see Figure 1). For explaining the OWL 2 DL reasoning we may be able to use built in functionality of e.g. the Pellet reasoner (cf. [14]).

A drawback of the current representation is that since conflicts are represented at the class level, they cannot be queried at the instance level. For instance, this means that we currently cannot query the system for all regions that have a land use which excludes that of a hypothetical new region: the exclusion relations do not hold between the `la:Land_Use` individuals directly, but are only inferred on the fly, for concrete situations. Ideally, one would therefore like the exclusions between types of land use specified at the class level, to propagate to all instances of these classes. One way to achieve this is by introducing a complex combination of OWL 2 DL role chains, self restrictions and the universal property [8, 15]. A similar approach has been described in [11] in the context of processes and actions, and we are currently investigating its use for the problems described here.

Future research will be directed towards more advanced spatial and numerical reasoning and the use of dynamic data to cater for the *planning* (compensate for the loss of x square meters of nature by creating the same area somewhere else) and *evaluation* scenarios specified in Section 2. A promising development is the recent addition of a spatial reasoning module to the SWI-Prolog engine in the POSEIDON project.¹⁹ Together with the Semantic Web library of this Prolog implementation, the reasoning services required by the application sketched here may be combined in one service.

¹⁹ See <http://www.swi-prolog.org> and <http://www.esi.nl/short/poseidon/> respectively

References

1. Boer, A., Hoekstra, R., Winkels, R., van Engers, T., Willaert, F.: *METAlex*: Legislation in XML. In Bench-Capon, T., Daskalopulu, A., Winkels, R., eds.: Jurix 2002: The Fifteenth Annual Conference. Frontiers in Artificial Intelligence and Applications, Amsterdam, IOS Press (2002) 1–10
2. Boer, A., Winkels, R., Vitali, F.: Metalex xml and the legal knowledge interchange format. In et al., P.C., ed.: Computable Models of the Law, LNCS 4884, Berlin, Springer Verlag (2008) 21–41
3. Winkels, R., Boer, A., Hupkes, E.: Legal atlas: Access to legal sources through maps. In: Proceedings of the 11th International Conference on Artificial Intelligence and Law, New York, ACM press (2007) 27–36
4. Boer, A., van Engers, T., Peters, R., Winkels, R.: Separating law from geography in gis-based egovernment services. *Artificial Intelligence & Law* **15**(1) (February 2007) 49–76
5. Winkels, R., Boer, A., Hoekstra, R.: CLIME: lessons learned in legal information serving. In van Harmelen, F., ed.: Proceedings of the European Conference on Artificial Intelligence (ECAI 2002), Amsterdam, IOS-Press (2002)
6. Valente, A.: Legal knowledge engineering: A modelling approach. PhD thesis, University of Amsterdam (1995)
7. van de Ven, S., Hoekstra, R., Breuker, J., Wortel, L., El-Ali, A.: Judging Amy: Automated legal assessment using OWL 2. In: Proceedings of OWL: Experiences and Directions (OWLED 2008 EU). (October 2008)
8. Motik, B., Patel-Schneider, P.F., Parsia, B., Bock, C., Fokoue, A., Haase, P., Hoekstra, R., Horrocks, I., Ruttenberg, A., Sattler, U., Smith, M.: OWL 2 web ontology language: Structural specification and functional-style syntax. Last call working draft, W3C (2008) (to be published, may be superseded).
9. Prakken, H.: Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law. Kluwer Law and Philosophy Library, Dordrecht (1997)
10. Doyle, J., Patil, R.: Two theses of knowledge representation: Language restrictions, taxonomic classifications, and the utility of representation services. *Artificial Intelligence* **48**(3) (1991) 261–298
11. Hoekstra, R.: Ontology Representation – Design Patterns and Ontologies that Make Sense. Volume 197 of Frontiers of Artificial Intelligence and Applications. IOS Press, Amsterdam (June 2009)
12. Hoekstra, R., Breuker, J.: Polishing diamonds in OWL2. In Gangemi, A., Euzenat, J., eds.: Proceedings of the 16th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2008). LNAI/LNCS, Springer Verlag (October 2008)
13. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. *Scientific American* **284**(5) (2001) 34–43
14. M. Horridge, B.P., Sattler, U.: Laconic and precise justifications in owl. In: Proc. of the 7th International Semantic Web Conference (ISWC 2008), Germany, Springer (2008)
15. Lutz, C., Sattler, U.: Mary likes all cats. In Baader, F., Sattler, U., eds.: Proceedings of the 2000 International Workshop on Description Logics, Aachen, Germany, CEUR-WS (August 17-19 2000)