



HAL
open science

Transformer-Based Argument Mining for Healthcare Applications

Valentina Leone, Luigi Di Caro, Serena Villata

► **To cite this version:**

Valentina Leone, Luigi Di Caro, Serena Villata. Transformer-Based Argument Mining for Healthcare Applications. *Artificial Intelligence and Law*, 2020, 28 (2), pp.207-235. <10.1007/s10506-019-09252-1>. <hal-03065217>

HAL Id: hal-03065217

<https://hal.science/hal-03065217v1>

Submitted on 15 Dec 2020

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

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



HAL Authorization

Taking Stock of Legal Ontologies: a Feature-based Comparative Analysis

Valentina Leone

University of Bologna, Bologna, Italy

University of Turin, Turin, Italy

University of Luxembourg, Esch-sur-Alzette, Luxembourg

Luigi Di Caro

University of Turin, Turin, Italy

Serena Villata

Université Côte d'Azur, CNRS, Inria, I3S, France

Abstract Ontologies represent the standard way to model the knowledge about specific domains. This holds also for the legal domain where several ontologies have been put forward to model specific kinds of legal knowledge. Both for standard users and for law scholars, it is often difficult to have an overall view on the existing alternatives, their main features and their interlinking with the other ontologies. To answer this need, in this paper, we address an analysis of the state-of-the-art in legal ontologies and we characterise them along with some distinctive features. This paper aims to guide generic users and law experts in selecting the legal ontology that better fits their needs and in understanding its specificity so that proper extensions to the selected model could be investigated.

Keywords Legal ontologies · Semantic Web · Modelling legal knowledge

1 Introduction

The modelling and the formalisation of legal knowledge are crucial aspects to implement in order to increase the automatic approach to the law field thus supporting the work of legal experts, enhancing legal information extraction and question answering systems and enabling automatic reasoning over legal cases.

Overlooking the first theoretical approaches to the formalisation of legal ontologies, such as the Functional Ontology of Law by Valente et al. (1994) or the frame-based ontology proposed by van Kralingen (1997), in the early 2000s most of the efforts focused on the modelling of core ontologies and knowledge

The authors have received funding from EU Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 690974 (MIREL).

interchange formats, such as LRI-Core by [Breuker and Hoekstra \(2004\)](#), CLO-Core Legal Ontology by [Gangemi et al. \(2005\)](#) and LKIF by [Hoekstra et al. \(2007\)](#).

Starting from the second decade of this century, the efforts concerning the legal knowledge representation moved towards the modelling of specific legal sub-fields as evidence of a greater awareness of the specificity which characterise each of them. This change of focus was accompanied by the consolidation of the Semantic Web as a reality for knowledge management and sharing. The Linked Data principles and the adoption of standardised knowledge representation formalisms as RDF and OWL are now common choices for publishing resources automatically accessible and processable through the Web.

However, despite the general acceptance of these good practices for the release of resources, the overall objective of a shared representation of legal knowledge has not been reached yet. The reuse of legal knowledge in fact requires a wide awareness of the already available resources which model the domain of interest. In order to evaluate a possible reuse, all the actors involved in the ontology building process, i.e., legal experts as well as developers, need to be constantly up-to-date about the state-of-the-art and they are expected to deepen the ontological commitment and the methodological choices adopted by each resource. If not, the risk is to create and release on the Web redundant representations of knowledge, which obstruct the economy of information promoted by the Semantic Web.

Considering that the last decade has seen a proliferation of ontologies and vocabularies which model different legal fields, it can be a good opportunity to take stock of the state-of-the-art concerning legal knowledge representation. Therefore, we propose a structured comparative analysis of the most recent legal ontologies and vocabularies. This work is mainly addressed to developers as well legal experts involved in the ontology building process. Our aim is to provide them with a practical source of information to consult in order to make an informed and conscious choice about the already modelled and reusable pieces of knowledge provided by other ontologies.

The paper is organised as follows: Section [2](#) describes the ontologies we analysed, Section [3](#) provides a description of the main features we used to study and classify them, and Section [4](#) discusses some insights resulting from our classification. In Section [5](#) future work directions end the paper.

2 Selected legal ontologies

In the past years, studies aiming at analysing and classifying legal ontologies have already been published. [Casellas \(2011\)](#) proposed a comprehensive survey about legal ontologies spanning a fifteen-years time range approximately, from early 90s to 2011. The ontologies features she considered in her analysis mainly concern the intended use of the ontology, the level of generality (core or domain), the degree of formalisation, the methodology used to build and evaluate the ontology, and its availability for reuse.




Fig. 1: The five domains according to which the analysed ontologies were grouped.

Recently, [de Oliveira Rodrigues et al. \(2019\)](#) enlarged the time-frame of their literature review and they analysed the legal ontologies proposed from late 90s to 2017. Their work presents different classification studies aimed at grouping ontologies among different dimensions, some of them similar to those already proposed by [Casellas \(2011\)](#). The new categorisation dimensions introduced by the authors concern the country and the venue where the literature about an ontology was published, its underlying legal theory, the syntactic and semantic peculiarities of legal texts that were addressed while producing the ontology (e.g., the dynamism of normative texts or the overlap of jurisdictions) and the legal subdomain it models.

If, on the one hand, the work of [Casellas \(2011\)](#) seems now out of date due to the lack of many recently developed ontologies, in [de Oliveira Rodrigues et al. \(2019\)](#) literature review it is difficult to identify the current emerging trends in the field due to the wide temporal interval their study focuses on. Moreover, information used to organise the ontologies in different types of classification were only collected from the scientific papers published to describe them. The ontologies documentation and the actual implementation, when available, seems not to have been taken into consideration. This methodology limits the analysis to a theoretical level which leaves out more technical details and deeper modelling choices.

Nowadays, the reuse of knowledge promoted by the Semantic Web principles require ontologists to exploit, as much as possible, the legal knowledge already made available through vocabularies, ontologies and knowledge graphs. To do so, experts who are involved in the ontology building task and who are planning to reuse an existing resource need to consider a wide set of details. Usually, those details are not limited to the theoretical features of an ontology, but also include more practical information, e.g. the on-line availability of the ontology source file or the presence of a specific class inside the ontology.

Starting from the Semantic Web principle of knowledge reuse, we take the classification of legal ontologies one step further by analysing the details of their implementation and including practical information concerning their actual availability for reuse. As both the aforementioned state-of-the-art literature reviews already analysed the resources produced in 90s and in the first decade of this century, we focused our attention on the most recently released legal ontologies. Thus, as an ideal continuation and extension of Casellas (2011) analysis, we considered the ontologies released from 2012, by the addition of two older ontologies which are still well known and used as it will be explained later. We excluded from our study the ontologies whose source files are not available for download. We make this decision to maintain consistency with our purpose to enable readers to analyse just the ontologies actually available to reuse. As it will be noticed in the following sections, only two ontologies do not accomplish this requirement. This is because they are very recent (less than two years old) and we believe that there is a possibility that they will be released later. Moreover, we decide to focus our attention on the resource that model a legal domain referring to some European or globally applicable legal framework. The ontologies that focus on a national jurisdiction were thus excluded from our analysis.

According to our selection criteria, we analysed a set of ten ontologies belonging to five domains related to different legal fields, as shown in Fig 1

1. *Policies*: it refers to the ontologies which model the permitted, mandatory and prohibited actions that can be made on a digital or material asset;
2. *Licences*: it includes the ontologies modelling the actions allowed on a resource protected by the intellectual property rights;
3. *Tenders and procurements*: this domain includes the ontologies modelling the processes used by public administrations and authorities to find contractors to entrust with services or supplies;
4. *Privacy*: the ontologies model the concepts concerning the protection of personal data.

Each domain is characterised by the different sources of law it refers to and by a distinctive jargon usually reflected in the classes and properties names of each related ontology.

In addition to the aforementioned domains, as showed in Fig 1 we analysed another set of four “cross-domains” ontologies which are difficult to associate to a specific legal field because they were proposed as a more generic model for expressing deontic operators (Normative Requirement Vocabulary), represent-

ing the content of legal texts in a machine-readable format (LegalRuleML) and indexing documents for search (Eurovoc and European Legislation Identifier).

In the following part of this Section, we provide a short description of each ontology.

2.1 Policies

2.1.1 Open Digital Rights Language

Open Digital Rights Language¹ (ODRL) is a language promoted by the ODRL Community Group² in order to model policies for digital content and media (Steyskal and Polleres (2014)). To do so, ODRL offers a Core Vocabulary to specify the minimum set of terms suitable to model the policies and a Common Vocabulary of general terms to model, for example, actions regulated by the obligations, permission and prohibitions expressed in the policies.

It models different types of policies, making a distinction between (i) a policy which is an agreement between an assigner and an assignee, (ii) a policy which is an offer from an assigner to an undefined wide audience and (iii) a policy which is a generic set of rules with no specified assigner and assignee.

Concerning the deontic logic, ODRL allows the expression of the effects associated to the non-compliance of an obligation, the effects of the non-compliance of some preliminary duties to obtain a permission and the duties to be accomplished for remedying to a violated prohibition. Finally, it is possible to associate a policy with some meta-information concerning, for example, its creator, its coverage (i.e., the jurisdiction applied upon the policy) and the reference to older versions of the policy.

2.1.2 Linked Data Rights ontology

The Linked Data Rights (LDR) ontology³ was developed by the Ontology Engineering Group⁴ and it is specifically designed to model the rights which can be exercised on a Linked Data resource. LDR ontology is based on ODRL from which it extends the classes *Action*, *Asset*, *Policy* and *Rule* in order to model the conditions of use of the Linked Data resources.

In detail, LDR defines three subsets of the ODRL *Action* class in order to represent the actions permitted on a resource protected by the intellectual property right, to use a database of Linked Data and to access a resource via the REST and SPARQL services. Moreover it defines which are the types of Linked Data resources (data-sets, link-sets, ontologies, resources and statements) and which are the types of policy that can be concluded (contract or licence).

¹ <https://www.w3.org/TR/odrl-vocab/>

² <https://www.w3.org/community/odrl/>

³ <http://oeg-dev.dia.fi.upm.es/licensius/static/ldr/>

⁴ <http://www.oeg-upm.net/>

As in this ontology there is also a reference to the intellectual property rights, but this is not the main focus, we included this ontology in the policy domain. However, it can be useful to take into account this ontology for the intellectual property field when the other models do not fit the needs of the users.

2.2 Licences

2.2.1 Creative Commons Rights Expression Language

The Creative Commons Rights Expression Language (ccREL⁵) is the standard promoted by Creative Commons⁶ (CC) to express the copyright licensing terms in a machine readable way. This ontology is more than six years old, but we decided to include it in this survey because of the wide dissemination of the Creative Commons licensing terms to regulate the use of resources protected by copyright.

The ccREL ontology models all the relevant actions provided by the Creative Commons standard, distinguishing among permissions, requirements and prohibitions. All of them are further specialised by the actions which allow the sharing of a work with third parties while maintaining the copyright. Moreover, the ontology allows the specification of the legal jurisdiction which applies on the modelled licence to be represented.

2.2.2 L4LOD

The Licence for Linked Open Data (L4LOD⁷) vocabulary uses a light ontological structure to organise the terms concerning licensing in the Web of Data. The deontic operators (permission, prohibition, obligation) are further specified in order to detail which actions can be necessarily or possibly made and avoided on Linked Open Data sources.

2.3 Tenders and procurements

2.3.1 LOTED2

LOTED2⁸ by Distinto et al. (2016), is a legal ontology which aims to represent the knowledge concerning the public procurements domain in the European Union. This ontology exploits the terminology contained in TED⁹, the reference online platform where all the public institutions of European

⁵ <https://www.w3.org/Submission/ccREL/>

⁶ <https://creativecommons.org/>

⁷ http://ns.inria.fr/l4lod/v2/l4lod_v2.html

⁸ <https://code.google.com/archive/p/loted2/source>

⁹ <https://ted.europa.eu/TED/main/HomePage.do>

and EEA countries publish their procurement notices. Starting from this website, LOTED2 enriches the TED lexicon with an ontological structure legally rooted on two European Union directives about the public contracts field: the Directive 2004/17/EC and the Directive 2004/18/EC. LOTED2 uses these two directives in order to model the legal concepts involved in the process of awarding a public contract, among which there are: the roles that an agent can play in the process, the different types of competition, the different types of documents used for the publication of a notice, the legal resources that regulate the field and the offers submitted for awarding a public contract.

The aforementioned aspects are all contained in the core version of LOTED2. An extended version of the ontology in which the concepts modelled in LOTED2 are integrated with some concepts and properties of the Good Relations is also available.

2.3.2 PPROC

The Public Procurement Ontology¹⁰ (PPROC), by Muñoz-Soro et al. (2016) aims to semantically represent the information published in official procurement documents, focusing on the Spanish law and in the EU law in general. Besides representing the usual information about tenders, PPROC objective is to represent the whole process of execution of tenders, starting from the publication of the contract until its termination.

Among its distinctive features, PPROC provides a classification of contracts according to different criteria, e.g. their administrative type or their subdivision in lots. Moreover it allows the specification of the criteria used for the evaluation of a tender, distinguishing them between subjective and objective criteria. The agents involved in a contract are expressed in the form of roles played during its execution and some hierarchies of roles are modelled. PPROC also represents the aspects which do not belong strictly to the set of properties of a tender or a contract, but which could be of interest for the suppliers (e.g., the kind of procedure followed during the execution of the procurement or its urgency).

It is important to remark that, in its attempt to model the public procurements and tenders domain, PPROC makes a big effort to try to reuse information already modelled in other existing ontologies, limiting the introduction of new classes and properties to very specific modelling requirements.

¹⁰ <http://contsem.unizar.es/def/sector-publico/pproc.html>

2.4 Privacy

2.4.1 Data Protection Ontology

The Data Protection Ontology^[11] by Bartolini et al. (2015) concerns the data protection field, as it is modelled in the GDPR (General Data Protection Regulation 2016/679). The Regulation came into force in May 2018, three years after the ontology published by Bartolini et al. (2015). However, even if the ontology is not based on the final version of the GDPR text, we decided to include this ontology to enable the interested reader to compare it with other two ontologies modelling the same field, that is GDPRtEXT (see Subsection 2.4.2) and PrOnto (see Subsection 2.4.4). This ontology is part of a more complex system where it plays the role of a knowledge base used to express data protection requirements as annotations inside a workflow model (e.g. a business process). The Data Protection Ontology was developed manually, extracting the terms of the domain of competence from a corpus of official normative sources. The main concepts modelled by the ontology concern the data protection principles, the rules of data processing and the rights of the data subject. In particular, the data protection principles are the glue that relates and justifies the duties of the data controller as well as the rights of the data subjects, making explicit the relation between a data subject right and the corresponding obligation for a data controller to guarantee this right.

2.4.2 GDPRtEXT

The GDPRtEXT^[12] (GDPR text extensions), by Pandit et al. (2018), is one of the most recent ontologies analysed in this survey and it deals with a currently central topic in the privacy domain: the aforementioned General Data Protection Regulation (GDPR).

The aim of GDPRtEXT is to represent the GDPR as a Linked Data resource, assigning an URI to each relevant part of the text. To do this, it extends some classes and properties of the ELI ontology (presented in Subsection 2.5.3) in order to specify the different parts in which the GDPR's text is structured (such as articles, recitals, citations and so on) and the properties that hold among them.

The ontology also provides more than 200 classes suitable to represent the relevant concepts introduced by the regulation and concerning the data protection field. The concepts' macro-areas modelled by the ontology are related to the categories of personal data, the concept of consent, the agents involved in the processing of the data, the actions that can be made on data, the rights

¹¹ <https://bitbucket.org/guerret/lu.uni.eclipse.bpmn2/src/3ca749d36cf193b9af8808c0fdf24858cdfeb21e/resources/dataprotection-rdf.owl?at=master&fileviewer=file-view-default>

¹² <https://openscience.adaptcentre.ie/ontologies/GDPRTtEXT/deliverables/docs/index-en.html>

of the data subject and the obligations of each agent which deals with the data.

GDPRtEXT also introduces a special property *isDefinedBy* which exploits the URI scheme created according to the Linked Data principles in order to link its classes to the relevant part of the text of the GDPR explaining the concepts they represent.

2.4.3 PrivOnto

PrivOnto is an ontology developed by [Oltramari et al. \(2018\)](#) in the context of the Usable Privacy Policy project¹³ and its aim is to model annotated privacy policies explaining the data practices implemented by a website.

PrivOnto was built from a corpus of 115 privacy policies of websites belonging to US-based companies. This corpus was annotated by some domain experts who were asked to identify the main categories representing data practices, together with their attributes. The result was a set of ten categories of data practices represented as frames. Each frame has its set of attributes together with the corresponding values, that refer to the fragment of the privacy policy they are taken from. Indeed, PrivOnto allows the modelling, with specific classes, of different parts of the text and the annotations associated to each of them.

As an application of this resource, a set of 57 different SPARQL queries was engineered in order to browse the annotated corpus over its different dimensions (categories, attributes and values).

2.4.4 PrOnto

Similarly to the Data Protection Ontology and GDPRtEXT (see Subsection 2.4.1 and Subsection 2.4.2), PrOnto (Privacy Ontology), proposed by [Palmirani et al. \(2018\)](#), focuses on the modelling of the knowledge concerning the GDPR. The purpose of PrOnto is not only to support information retrieval, but also to provide a model on which techniques of legal reasoning and compliance checking could be applied.

Among its distinctive features, PrOnto focuses on the distinction between agents and roles, with the former able to cover particular roles inside different contexts and for a limited interval of time. Moreover, PrOnto models the sequence of actions aimed at processing personal data. Specifically, it makes a distinction between a planned sequence of actions named *workflow* and the real execution of this plan, named *workflow execution*. A temporal reference can be associated to each action and some boolean attributes are associated to the workflow in order to represent and automatically infer its lawfulness, fairness and transparency.

¹³ <https://www.usableprivacy.org/>

Besides the traditional deontic operators, (i.e., permissions, prohibitions, obligations and duties) PrOnto explicitly models compliance with and violation of an obligation by relating the *obligation* class with the *compliance* and *violation* classes as well as a *right* with the corresponding *permission*.

Within the DAPRECO project by Bartolini et al. (2016), the PrOnto ontology has been associated to fine-grained if-then rules in reified Input/Output logic (Robaldo and Sun (2017)). Rules represent GDPR norms and are encoded in LegalRuleML (see Section 2.5.2). To date, this the biggest knowledge base in LegalRuleML freely available online¹⁴

2.5 Cross-domains ontologies

2.5.1 Eurovoc

Eurovoc¹⁵ is a multilingual and multidisciplinary thesaurus managed by the Publications Office of the European Union. Its function is to index the documents issued by the European Union Institutions in order to ease their retrieval.

The concepts are organised in 21 sectors which in turn are composed by micro-thesauri. Each sector concerns a field of competence of the European Union and each concept can be associated with only one sector to avoid ambiguities (except for the sector *Geography* which allows a polihierarchy).

Each concept is lexicalised by a set of terms in which only one is the *preferred term* (i.e., the term used for the indexing of the concept), while the others are the *non preferred terms* (i.e., synonyms of the preferred term not used for the indexing of the concept they represent). All the terms associated to a concept are provided with their translations in all the 23 languages spoken inside the European Union and Macedonian, Serbian and Albanian. Nevertheless, while there is a unique correspondence between the different translations of a preferred term, the set of the non preferred terms associated to a concept can vary considering their representation in different languages in order to maintain the linguistic nuances of each national legal lexicon.

The terms in Eurovoc are also linked to each other through some semantic relations: beside the classical hierarchical one, also associative relations can be found among terms that are semantically related but are not on the same hierarchical structure.

Although the project which led to the creation of Eurovoc is more than twenty years old, its updating is constant and frequent: the thesaurus is continuously enriched with new terms concerning the topics dealt by the EU and cleaned up by removing obsolete terms.

¹⁴ https://github.com/dapreco/daprecokb/blob/master/gdpr/rioKB_GDPR.xml

¹⁵ <https://publications.europa.eu/en/web/eu-vocabularies/th-dataset/-/resource/dataset/eurovoc>

2.5.2 LegalRuleML

LegalRuleML¹⁶ by Palmirani et al. (2011) and Athan et al. (2015), is a project promoted by the OASIS LegalRuleML Technical Committee¹⁷ which aims to develop a standard for the legal knowledge representation and exchange. To reach this goal, LegalRuleML offers a markup language which permits the harmonisation of different types of legal texts, such as norms, guidelines and policies.

Even though LegalRuleML is not properly an ontology but a markup language, we decided to include this resource inside our survey because it provides a rich set of concepts and properties which enable the management of the complexities of a formal representation of legal texts in a machine-readable way. Among its distinctive features, LegalRuleML provides some parameters to model the different interpretations that could be associated to a rule, to keep track of the author of a document or its fragments, to manage the temporal evolution of the norms and to take into account the defeasibility of the law.

Thus, the advantage and the final goal of LegalRuleML is the possibility to maintain the same expressive power independently from the way the norm is expressed, using the natural language or a formal machine-readable representation.

2.5.3 European Legislation Identifier ontology

The European Legislation Identifier (ELI) ontology¹⁸ is a model which allows the publication of legal documents of different European Union countries using a shared and uniform set of metadata in order to enhance interoperability among the national administrations. Nowadays, this resource is used by 11 of the 28 EU countries and by the EU Publication Office.

According to the information published by ELI Task Force (2018), the ELI ontology reflects many of the basic principles of FRBR (Functional Requirements for Bibliographic Records) vocabulary¹⁹, contextualising them into the legal field. While the FRBR provides the description of a bibliographic record in terms of *work*, *expression*, *manifestation* and *item*, the ELI ontology describes a legal document through the concepts of *legal resource*, *legal expression* and *format*. In detail, *legal resource* refers to the intellectual creation, independently from its translation in more than one language and from the format used for its publishing; it corresponds to the *work* property in FRBR. The *legal expression* concept is the realisation of a *legal resource* using a sequence of signs as, for examples, the alphanumeric characters and it corresponds to the *expression* property in FRBR. The *format* refers to the physical means

¹⁶ <http://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/legalruleml-core-spec-v1.0.html>

¹⁷ https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=legalruleml

¹⁸ <https://publications.europa.eu/en/web/eu-vocabularies/model/-/resource/dataset/eli>

¹⁹ <https://sparontologies.github.io/frbr/current/frbr.html>

used to store the *legal expression* (could be paper or an electronic format) and it corresponds to the *manifestation* property on FRBR. However, the *item* property of FRBR does not have a correspondence in the ELI ontology.

Since the documents issued by different EU countries could be described with different metadata according to the national jurisdiction they refer to, the ELI ontology overlooks these differences in order to represent only the common metadata of the national legal documents, providing the user the possibility to personalise and extend the set of metadata according to its needs. Therefore, the set of properties that can be established among the aforementioned three classes is not so large and they mainly concern the type of the represented document, the topics it deals with, the entry into force and the legal value of the document according to the format it is represented with.

2.5.4 Normative Requirements Vocabulary

The Normative Requirements Vocabulary²⁰ (NRV), by Gandon et al. (2017), is an ontology which extends LegalRuleML and whose aim is to exploit the standard frameworks offered by the Semantic Web in order to represent normative requirements and rules. Differently from other existing legal ontologies, NRV is not limited to the representation of the three main deontic operators (i.e. permission, obligation and prohibition), but it specifies and organises them in a hierarchical structure according to different criteria which concern: the need for compensation, the possibility to breach or fulfil a requirement and the temporal aspects involved in their validity and compliance.

NRV also uses the named graphs of RDF 1.1 in order to represent the states of affairs, that is the contexts on which the deontic operators can be applied. Then, given that OWL does not support the named graph structure, a SPARQL approach is tested for making complex inferences in which the formalised normative requirements are applied upon a state of affairs.

3 Features Description

This Section contains a description of each feature we used to classify the legal ontologies. We organised the overall set of features in three macro-classes according to the type of property modelled by the features they include. More specifically, we distinguish between:

- *general information* class: it contains several features about the ontology disclosure and the purpose of its creation;
- *modelling information* class: it refers to the methodological and technological choices followed in order to build the ontology;
- *semantic information* class: it groups all the features concerning the way in which the ontology models the knowledge it refers to.

²⁰ http://ns.inria.fr/nrv/v1/nrv_v1.html

Table 1: The macro-classes and the corresponding features used to classify the legal ontologies

Macro-class	Features
general information	extended name legal domain purpose year current version licence updates frequency number of references link
modelling information	development construction language knowledge sources for terms extraction external vocabularies references ground ontology level of structure knowledge representation formalism axioms design patterns evaluation
semantic information	modelling of temporal aspects adopted normative model deontic logic model

As mentioned before, each of these macro-classes is a set of more specific features as detailed in Table 1. In the following part of the Section, we provide a description of each feature used to classify the analysed legal ontologies.

3.1 General information class

As mentioned above, the features contained in this class refer to the generic purpose for which the ontology was built together with some practical information useful for those who are actually interested in using the resource. Eight features belong to this class.

The first information concerns the *extended name* of the ontologies. As they are often referenced by their acronyms in literature, their full name could provide to the reader a first insight of the scope of the ontology, also helping her to memorise the acronym itself.

The *legal domain* feature refers to one of the five domains listed in Section 2 and it corresponds to the visual information represented in Fig 1. This feature is further specified by *purpose* which contains a brief description of the main scope and function of the ontology inside the specified domain. Finally, the *year* feature indicates the year of the ontology first release.

Together with this general information, we decided to include some more specific features in order to provide the readers with useful information con-

cerning the retrieval of an ontology on the Web and its reuse. To this purpose, the *current version* feature refers to the most recent released version of the ontology, while *licence* provides the information concerning the licence under which a resource is made available for reuse. Such feature could help interested users to fairly use the ontology, respecting any limitation and constraint in its adoption. Then, to assess the frequency of updates made to an ontology, we introduced the *updates frequency* feature, whose possible values are: *low*, *medium* and *high*. In the following tables, the date of the last update is provided in brackets. This feature is important to understand if the resource already reached a stable point and to evaluate if it is kept up-to-date according to the changes of the domain that it models.

In order to provide readers with an estimate of how much an ontology is known, we also include the feature *number of references*. To estimate this number, we used the Google Scholar²¹ search engine and, for each paper describing an ontology and included in the bibliography of this study, we took the number of references from its publication date until May 2019, as returned by Google Scholar. For resources which do not have a reference paper, we searched from the number of citation starting from 2012 in order to be consistent with the year we chose to start our ontology collection (see Section 2). Moreover, we used to research keywords: the first one contained the extended name of the ontology followed by the term “ontology” (except for Eurovoc, where we used the term “thesaurus” as it is usually associated to this resource), while the second one contained the corresponding acronym (if available) followed again by the term “ontology”. The two keywords were then linked by a disjunction operator (i.e., OR). For instance, for the ELI ontology we built the following string: “*European Legislation Identifier ontology*” OR “*ELI ontology*”, where the quote marks were used to obtain only exact matches.

Finally, the *link* feature specifies the at-present active link to the Web page containing the ontology documentation. Usually, if available, this Web page also contains the link to download the ontology source file.

Table 2, Table 3 and Table 4 classify the ontologies presented in Section 2 according to these features.

²¹ <https://scholar.google.it/>

	Eurovoc	ccREL	LegalRuleML	ODRL	L4LOD
Extended name	European Vocabulary	Creative Commons Rights Expression Language	Legal Rule Modeling Language	Open Digital Rights Language	Licenses for Linked Open Data
Legal domain	cross-domains	licences	cross-domains	policies	licences
Purpose	indexing of the documentary information of the EU institutions	machine-readable standard to express licensing terms	modelling of legal norms allowing legal reasoning	representation of the conditions of usage of digital assets	representation of existing licensing terms in the Web of Data
Year	1984	2008	2011	2012	2013
Current version	4.9	unique version	1.0	2.2	0.2
Licence	commercial or non-commercial use allowed providing appropriate acknowledgement	CC BY 3.0	OASIS Intellectual Property Rights Policy	W3C Community Contributor License Agreement (CLA)	CC-BY-SA
Updates frequency	high (29 Mar 2019)	low (1 May 2008)	high (8 May 2018)	high (15 Feb 2018)	low (10 May 2013)
Number of references	280	190	56	21	31
Link	bit.ly/2MY0TpM	bit.ly/2Lua4gp	bit.ly/2sxpskV	bit.ly/2J75JPj	bit.ly/2m40FSn

Table 2: Classification of ontologies published from 1984 to 2013 according to the *general information* class of features. Last revision of the information contained in the table: May 2019.

	ELI	LOTED2	PPROC	LDR	Data Protection Ontology
Extended name	European Legislation Identifier	not found	Public Procurement Ontology	Linked Data Rights ontology	not applicable
Legal domain	cross-domains	tenders and procurements	tenders and procurements	policies	privacy
Purpose	metadata for the description of legal documents issued by the EU and its Member States	indexing, search and retrieval of European public procurement notices	management of public procurements and the execution of contracts	representation of policies of Linked Data resources	Model the GDPR concepts, focusing on the obligation of the data controller in relation with the rights of the data subject
Year	2014	2014	2014	2014	2015
Current version	1.2	unique version	1.0.0	unique version	unique version
Licence	reuse allowed providing appropriate acknowledgement	GNU GPL v3	CC BY-SA 4.0	CC BY 4.0	not found
Updates frequency	low (21 Nov 2018)	low (16 Jan 2014)	low (29 Oct 2014)	low (1 Sep 2014)	low (16 Feb 2016)
Number of references	27	23	10	4	10
Link	bit.ly/2NyUimC	bit.ly/2m5os4q	bit.ly/2MwxGpQ	bit.ly/2KU59cx	bit.ly/2uhumDv

Table 3: Classification of ontologies published from 2014 to 2015 according to the *general information* class of features. Last revision of the information contained in the table: May 2019.

	GDPRtEXT	NRV	PrivOnto	PrOnto
Extended name	GDPR text extensions	Normative Requirements Vocabulary	not found	Privacy Ontology
Legal domain	privacy	cross-domains	privacy	privacy
Purpose	representation of the GDPR concepts with a direct link to the Regulation text	representation of annotated privacy policies of websites	representation of normative requirements and rules as LOD	representation of the GDPR concepts for legal reasoning and compliance checking
Year	2017	2017	2017	2018
Current version	0.6	unique version	unique version	not yet released
Licence	CC by 4.0	not found	not found	not applicable
Updates frequency	high (31 July 2018)	law (last update not found)	not found	not applicable
Number of references	10	6	20	5
Link	bit.ly/2xwjTZJ	bit.ly/2KFwkIC	not found	not applicable

Table 4: Classification of ontologies published from 2017 to 2018 according to the *general information* class of features. Last revision of the information contained in the table: May 2019.

3.2 Modelling information class

The eleven features contained in this class concern all the modelling choices which are immediately reflected in methodologies and standards used to build the ontologies.

The *language* feature refers to the main natural language used to specify the concepts, the relations and the lexicon inside the ontology while *development* indicates the approach adopted in the ontology building process, that is a bottom-up approach (from lexicon to concepts), a top-down approach (from legal foundations to lexicon) or a middle-out approach, which merges the techniques of the previous two methods.

The *construction* feature specifies if the modelling of the ontologies' concepts and relations was manual or used some Natural Language Processing (NLP) technique to partially automatise the process of building the ontology. Linked to this aspect, two features concern the sources from which the concepts inserted in the ontology were chosen. The first one is *knowledge source (KS) for terms extraction*, that is legal documents or websites used to extract the relevant concepts and the corresponding ontology lexicon. In contrast, the *external vocabulary (EV) reference* feature refers to the existing ontologies and vocabularies which the ontology reuses specifying the URIs of some of their concepts and properties. Therefore, the difference between these two last features is that the legal documents listed in correspondence of the first feature only provide the raw concepts which are relevant for the domain but which needed to be formally modelled before being inserted in the ontology, while the second feature looks at the reuse of some parts of existing ontologies in order to adopt some concepts and relations already modelled by them. Similarly, the *ground ontology* feature refers to the main ontology which is extended by the analysed resource. This feature can be seen as a specialisation of *external vocabulary reference*. The difference is that an ontology which uses another one as ground ontology inherits from it the great part of its concepts and structure, while an ontology that makes some reference to external vocabularies adopts its own structure and reuses only some concepts of other existing resources.

The *level of structure* feature is a quantitative evaluation of the number of concepts and relations modelled by the ontology. This property can be expressed by three values that denote a growing number of classes and relations: *lightly structured*, *moderately structured* and *highly structured*. The *knowledge representation (KR) formalism* refers to the formal language used to represent the ontology in a machine readable way. At present, the two *de facto* standards used to represent ontologies are RDF and OWL. Connected to this feature, the *axioms* feature is also considered. The feature refers to the three possible level of axioms planned by OWL 2 specification: class expression axioms, object property axioms and data property axioms.

Taking into account the principle of reuse promoted by the Semantic Web, we also considered the *ontology design patterns* used to represent some parts of knowledge whose modelling was already codified in a standard representa-

tion. Finally, the *evaluation* feature analyses which methods were adopted to evaluate the created knowledge model provided by the ontology.

Tables from [5](#) to [7](#) classify the analysed ontologies according to the features of this class.

	Eurovoc	ccREL	LegalRuleML	ODRL	L4LOD
Development	not found	not found	not found	not found	not found
Construction	manual	manual	manual	manual	manual
Language	EUs languages, Macedonian, Albanian, Serbian	English	English	English	English
KS for terms extraction	ECLAS thesaurus, SCAD, EC-01, Official Gazette indices	not found	not applicable	not applicable	not found
EV references	FRBR, Dublin Core, SKOS	not found	not found	Dublin Core, SKOS, FOAF	not found
Ground ontology	none	none	RuleML	none	none
Level of structure	lightly structured	lightly structured	highly structured	highly structured	lightly structured
KR formalism	RDF	RDF	RelaxNG and XML Schema, RDFS, XSLT	RDF	RDF
Axioms	not found	class level, property level	class level, property level	class level, property level	class level
Design patterns	not found	not found	container, collection, recursive element, marker, composite	not found	not found
Evaluation	EU institutions, Publication Office, national and regional parliaments	not found	not found	not found	not found

Table 5: Classification of ontologies published from 1984 to 2013 according to the *modelling information* class of features. Last revision of the information contained in the table: May 2019.

	ELI	LOTED2	PPROC	LDR	Data Protection Ontology
Development	not found	middle-out	bottom-up	not found	bottom-up
Construction	manual	manual	manual	manual	manual
Language	English	English	English	English	English
KS for terms extraction	not applicable	TED website, EI Directive 2004/17/EC and EU Directive 2004/17/EC	buyer profiles, EU directives, public procurements' announcement models of Spanish legislation	not found	GDPR, Data Protection Directive (DPD), Handbook on European data protection law
EV references	FRBR, Dublin Core, SKOS	LKIF-core, GoodRelations	CPV, PCO, FOAF, SKOS, DC, Organization Ontology, schema.org, Good Relations	ODRL, SKOS	LKIF-Core, SKOS
Ground ontology	FRBR/RDA	none	none	ODRL	none
Level of structure	lightly structured	moderately structured	highly structured	lightly structured	lightly structured
KR formalism	OWL	OWL	OWL	OWL	OWL
Axioms	class level	class level	class level	class level, property level	class level
Design patterns	not found	Social Reality	not found	not found	not found
Evaluation	provided by users	not found	provided by two Spanish public authorities	not found	not found

Table 6: Classification of ontologies published from 2014 to 2015 according to the *modelling information* class of features. Last revision of the information contained in the table: May 2019.

	GDPRtEXT	NRV	PrivOnto	PrOnto
Development	bottom-up	not found	middle-out	bottom-up following the MeLOn methodology
Construction	manual	manual	manual	manual
Language	English	English	English	English
KS for terms extraction	GDPR, document issued by official sources, industry-based sources	not applicable	115 privacy policies of US-based companies	GDPR, terms of use, information, privacy policies, consent forms
EV references	ELI ontology	LegalRuleML, RuleML	not found	ALLOT, FRBR, LKIF Core, PWO, LegalRuleML metamodel
Ground ontology	ELI ontology	LegalRuleML	none	none
Level of structure	lightly structured	moderately structured	lightly structured	highly structured
KR formalism	OWL	RDF	OWL	OWL
Axioms	class level	class property	class level	class level
Design patterns	not found	not found	not found	Time-indexed Value in Context, Time interval
Evaluation	not found	SPARQL queries	not found	SPARQL queries

Table 7: Classification of ontologies published from 2017 to 2018 according to the *modelling information* class of features. Last revision of the information contained in the table: May 2019.

3.3 *Semantic information class*

So far, we presented a set of features which are independent from the legal domain and which could be applied potentially to analyse and compare the ontologies belonging to every domain of interest. In this Section, we analyse three features which specifically refer to the way in which the legal knowledge is modelled.

The *modelling of temporal aspects* feature allows to specify if an ontology models some temporal aspects concerning the legal field of interest and provides a brief description of the way in which this is done. There are a lot of different possibilities to model a temporal feature inside an ontology: it could be a simple time mark associated to the issue of a policy, or an interval of time which specifies the validity of an obligation or, again, it could be an implicit representation of time which focuses on the parameters that could vary over it, e.g., the status of a norm or the jurisdiction under which it is valid.

When an ontology permits the modelling of norms and rules, the *adopted normative model* feature specifies the type of rules that the ontology can represent (e.g., constitutive rules, prescriptive norms, etc.). Finally, the *deontic logic model* feature provides a short description of the deontic operators modelled inside the ontology (i.e., obligation, duties, permissions and rights). As for the previous feature, this one holds only if the ontology deals with norms and rules. However, since norms are one of the main focus of the legal domain, a lot of the analysed ontologies model the deontic operators. For example, some of them only represent permissions, obligations and prohibitions, others model also the violations of obligations and prohibitions, while others provide a hierarchy of deontic operators organising them according to different criteria (e.g., temporal criteria or need for compensation of a violated norm).

The classification of the analysed ontologies according to these three features is provided in Tables from [8](#) to [10](#).

	Eurovoc	ccREL	LegalRuleML	ODRL	L4LOD
Modelling of temporal aspects	not applicable	not found	modelling of the aspects of a rule that vary over time (e.g., status, validity, jurisdiction)	modelling of the date and time a policy is issued or modified. Date and time constraint on the validity of a deontic operator.	not found
Adopted normative model	not applicable	prescriptive rules	constitutive, technical and prescriptive rules	prescriptive rules	prescriptive rules
Deontic logic model	not applicable	requirements and prohibitions set by the Creative Commons standard	permission, rights, obligation, prohibition, compliance with a prohibition or an obligation, violation of a prohibition or an obligation, reparation of a violation	permissions, prohibitions and obligations over a digital or material asset	permissions, obligations and prohibitions over the licensed data

Table 8: Classification of ontologies published from 1984 to 2013 according to the *semantic information* class of features. Last revision of the information contained in the table: May 2019.

	ELI	LOTED2	PPROC	LDR	Data Protection Ontology
Modelling of temporal aspects	not applicable	date and time associated to tenders	only to indicate the deadline for submissions of tenders and requests of participation	not found	not found
Adopted normative model	not applicable	not applicable	not applicable	prescriptive rules	prescriptive rules
Deontic logic model	not applicable	not applicable	additional obligations that a contract needs, requirements that a tender needs in order to be submitted	right over a Linked Data resource	obligation (of the data controller) and rights (of the data subject)

Table 9: Classification of ontologies published from 2014 to 2015 according to the *semantic information* class of features. Last revision of the information contained in the table: May 2019.

	GDPRtEXT	NRV	PrivOnto	PrOnto
Modelling of temporal aspects	information about personal data retention and storage period are modelled as ontology classes	temporal aspects are modelled through the concepts of perdurance, persistence, co-occurrence and preemptiveness of a deontic operator	not applicable	temporal intervals associated to actions in workflows, to agents' roles and to deontic operators
Adopted norms model	prescriptive rules	prescriptive rules	prescriptive rules	constitutive and prescriptive rules
Deontic logic model	obligations of different agents mentioned in the GDPR (Controller, Processor, Data Protection Officer) and right of the data subject	permissions, obligations and prohibition are organised according to the principles of compensation, compliance, violation, temporal validity and realisation	not applicable	permissions, prohibitions, obligations, rights, compliance with an obligation and violation of an obligation. Some references are modelled between obligations and compliance/violation and between rights and permissions

Table 10: Classification of ontologies published from 2017 to 2018 according to the *semantic information* class of features. Last revision of the information contained in the table: May 2019.

4 Concluding remarks

The analysis of the ontologies contained in this survey and the completion of the tables included in the previous Section led us to a greater awareness about some weaknesses concerning the panorama of the existing legal ontologies. The remarks we made can be grouped according to the division in macro-classes used to organise the features described previously.

Concerning the general information about an ontology (summarised in the *general information* class) some lack of standardisation still exists in the graphical user interfaces (GUIs) used to make the ontology scope and content available to the final user. Currently, the LODE²² tool is one of the most common Web services used to automatically create these GUIs. LODE processes the *owl* file of an ontology to create an HTML page which lists classes, properties and axioms of the ontology together with some metadata indicating the author(s), the release date, the current version and the licence of the ontology, as shown in Fig.2

An unified look for the GUIs exposing the content of an ontology could be helpful for users concerned with ontology building and reuse, as it could reduce the time spent to look for the information within websites.

Linked to this problem, the second issue concerns the need to make explicit all the details concerning the download and the licence of an ontology. Browsing the Web pages of the different ontologies, it was sometimes difficult for us to find this information. However, it seems clear that without them, a fair reuse of the ontologies would not be promoted.

A special case concerns the resources made available by the European Union whose orientation towards the Semantic Web and the Linked Open Data is remarkable. They are all collected in the EU vocabularies portal²³ where a tab-like GUI organises all the information about a resource as it shown in Fig.3

As it can be noted, this interface is very different from the GUI which can be created with LODE. Even if the download links are well visible, the type of licence which regulates the use of each resource is not specified. We found this information in the old Web sites of each resource, before their grouping inside the portal, under the heading “Legal notice”. Moreover, in the current interface of the EU vocabularies portal, the title of each tab sometimes does not clarify the information associated with it, and the documentation of the different resources is not standardized. For example, the documentation of ELI is an *xlsx* file which must be downloaded and opened with a commercial software in order to be visualized. In contrast, the description of Eurovoc is better organized into expandable windows inside the tab.

Therefore, according to these remarks, some improvement would be desirable to harmonize the way in which the metadata on legal ontologies issued by the EU are organised inside the portal.

²² github.com/essepuntato/LODE

²³ publications.europa.eu/en/web/eu-vocabularies



Normative Requirements Vocabulary

IRI:

<http://ns.inria.fr/nrv#>

Authors:

Fabien Gandon

<http://ns.inria.fr/fabien.gandon/foaf.html#me>

Contributors:

Fabien Gandon

Guido Governatori

Publisher:

Inria

Other visualisation:

[Turtle/N3](#), [RDF/XML](#), [Onto-Doc](#)

Abstract

An OWL vocabulary for describing normative requirements.

Table of Content

1. [Classes](#)
2. [Object Properties](#)
3. [Named Individuals](#)
4. [Namespace Declarations](#)

Classes

achievement	atomic formula	co-occurant	punctual requirement	compensable requirement	compensat
maintenance	modality	non co-occurant	punctual requirement	non compensable requirement	non perc
Normative Requirement	obligation	penalty statement	perdurant achievement requirement	permission	
suborder list	term	violable requirement	violated requirement	violation	

Fig. 2: An excerpt of the NRV GUI, automatically generated using the LODE tool.

Concerning the methodological and technological choices made during the development of an ontology, this information is never displayed on the aforementioned GUIs and it could be difficult to find also reading the literature published together with the ontology. However, this information is important for several reasons: first of all, it provides a scientific foundation to the work allowing other researchers to analyse and verify it, secondly, it enables an easy and understandable interpretation of the corresponding literature in which this information is sometimes implicit, even if it is at the basis of the development of the ontology.

A positive aspect that we noticed during the analysis of the proposed resources is the trend promoted by the Semantic Web principles to reuse the concepts and the properties of other ontologies or to propose extensions of existing ontologies using them as *ground ontologies*. However, we noticed a lack of sensitivity to the adoption of the ontology design patterns (ODPs) in

Asset

European Legislation Identifier (ELI)

Ontology

URI: <http://publications.europa.eu/resource/dataset/eli>

About Documentation

The European Legislation Identifier rests on three pillars:

1. Identification of legislation: URI templates at the European, national and regional levels based on a defined set of components
2. Properties describing each legislative act: Definition of a set of metadata and its expression in a formal ontology
3. Serialisation of ELI metadata elements: Integration of metadata into the legislative websites using RDFa

ID: <http://publications.europa.eu/resource/eli>

Version: V1.2 **LATEST**

Published: 2018-11-21

Author: [Publications Office](#)

Publisher: [Publications Office](#)

Downloads

[eli-20181121-0.zip](#)

[eli-sdo-mapping.xlsx](#)

[eli-sdo.ttl](#)

[eli.owl](#)

[eli_ontology.xlsx](#)

Other versions

V1.1

V1.0

Fig. 3: Some information about ELI as displayed on the EU vocabularies portal.

the ontology building process. As outlined by [Gangemi and Presutti \(2009\)](#), ODPs are modelling solutions to solve recurrent ontology design problems. The ODPs differ from the reuse of single concepts as they are micro-ontologies which model a piece of knowledge which occurs frequently in different domains. The low use of ontology patterns could be associated to the difficulty to identify, inside a complex modelling problem, the parts which could be covered by an ODP because it requires the knowledge of the full landscape of available ODPs. However, some portals ease their retrieval collecting the existing design patterns (among them we mention www.gong.manchester.ac.uk/odp/html and www.ontologydesignpatterns.org).

Finally, the most important lack that we noticed in the features involving the *modelling information* class is evaluation. In the literature related to the resources, we have not often found any mention to the criteria used to evaluate the proposed models. However, as shown in [Table 7](#) the current trend is to provide SPARQL queries to test the validity of some competencies questions and the fulfilment of some objectives which the ontology should reach. This is especially done by the most recent ontologies as for example NRV and PrOnto. In contrast, older ontologies mention in their literature the fact that they are used by real users, as in the case of PPROC or the resources released by the European Union. We can consider it as a method of evaluation since the actual use of a resource is one of the best ways to test the robustness of a knowledge model.

The considerations we made concerning the *semantic information* class call back the aforementioned problem of the ontologies design patterns. Indeed, we noticed that each ontology models a specific legal domain and adopts its own ontological commitment, with a consequent proliferation of different knowledge models referring to similar use cases. For example, the deontic operators, being one of the main focus of different legal domains, are modelled in many ontologies but the aspects that each of them considers are different. For example, some ontologies associate a temporal reference to the validity of an operator (as LegalRuleML or ODRL do) while others do not (e.g. L4LOD). Or, again, some ontologies make a distinction between an obligation which is respected and an obligation which is violated (as NRV), while others not (e.g. LDR). Thus, even if the legal domain has plenty of recurrent use cases, few efforts are dedicated to find a standardized solution to design problems which recur often within the legal domain.

5 Future perspectives

According to the remarks proposed in the previous Section, some improvements could be done to enhance an ontology building process oriented towards the reuse of existing resources.

First of all, the creation of a new set of metadata to include inside the ontology source file should be evaluated in order to complete the information that is already showed in the graphical interfaces displaying the content of an ontology. We believe that the most needed information is both of a general and of a legal nature. In the first instance, some metadata for indicating the methodology of development followed to create the ontology and the embedded design patterns would be useful to ensure the reuse of the ontology itself. In the second instance, we think about a set of metadata able to summarise some of the purely legal aspects modelled into an ontology. Some of these metadata could recall some of the features used inside this survey to classify the ontologies, as for example the modelled deontic operators and the type of modelled norms (if this feature is applicable).

In addition to a new set of metadata for the description of the ontology features, it could be important to address the problem pointed out at the end of Section 4 concerning the need of legal design patterns to reuse inside the ontologies. Some witnesses in this direction are provided by [Haapio and Hagan \(2016\)](#) and [Haapio et al. \(2018\)](#). An effort to discover recurrent legal knowledge and to model it in the form of a standardises legal use case with the corresponding ontology design pattern could improve the quality of the released ontologies reducing the efforts required to model legal knowledge. This is especially true considering that usually the design of ontology-based systems is assigned to computer scientists who need, in addition to the technical background, a further knowledge about the legal domain which usually they do not hold.

References

- Tara Athan, Guido Governatori, Monica Palmirani, Adrian Paschke, and Adam Wyner. Legalruleml: Design principles and foundations. In *Reasoning Web International Summer School*, pages 151–188. Springer, 2015.
- Cesare Bartolini, Robert Muthuri, and Cristiana Santos. Using ontologies to model data protection requirements in workflows. In *JSAI International Symposium on Artificial Intelligence*, pages 233–248. Springer, 2015.
- Cesare Bartolini, Andra Giurgiu, Gabriele Lenzi, and Livio Robaldo. Towards legal compliance by correlating standards and laws with a semi-automated methodology. In *Benelux Conference on Artificial Intelligence*, pages 47–62. Springer, 2016.
- Joost Breuker and Rinke Hoekstra. Epistemology and ontology in core ontologies: Folaw and Iri-core, two core ontologies for law. In *In Proceedings of the EKAW04 Workshop on Core Ontologies in Ontology Engineering*, pages 15–27. Northamptonshire, UK, 2004.
- Núria Casellas. *Legal ontology engineering: Methodologies, modelling trends, and the ontology of professional judicial knowledge*, volume 3. Springer Science & Business Media, 2011.
- Cleyton Mário de Oliveira Rodrigues, Frederico Luiz Gonçalves de Freitas, Emanuel Francisco Spósito Barreiros, Ryan Ribeiro de Azevedo, and Adauto Trigueiro de Almeida Filho. Legal ontologies over time: A systematic mapping study. *Expert Systems with Applications*, 130:12–30, 2019.
- Isabella Distinto, Mathieu d’Aquin, and Enrico Motta. Loted2: An ontology of european public procurement notices. *Semantic Web*, 7(3):267–293, 2016.
- ELI Task Force. *ELI Implementation Methodology: Good Practices and Guidelines*. Publications Office, 2018.
- Fabien Gandon, Guido Governatori, and Serena Villata. Normative requirements as linked data. In *The 30th international conference on Legal Knowledge and Information Systems (JURIX 2017)*, 2017.
- Aldo Gangemi and Valentina Presutti. Ontology design patterns. In *Handbook on ontologies*, pages 221–243. Springer, 2009.
- Aldo Gangemi, Maria-Teresa Sagri, and Daniela Tiscornia. A constructive framework for legal ontologies. In *Law and the semantic web*, pages 97–124. Springer, 2005.
- Helena Haapio and Margaret Hagan. Design patterns for contracts. In *Networks. Proceedings of the 19th International Legal Informatics Symposium IRIS*, pages 381–388, 2016.
- Helena Haapio, Margaret Hagan, Monica Palmirani, and Arianna Rossi. Legal design patterns for privacy. In *Data Protection/LegalTech. Proceedings of the 21th International Legal Informatics Symposium IRIS*, pages 445–450, 2018.
- Rinke Hoekstra, Joost Breuker, Marcello Di Bello, Alexander Boer, et al. The IRI core ontology of basic legal concepts. *LOAIT*, 321:43–63, 2007.
- José Félix Muñoz-Soro, Guillermo Esteban, Oscar Corcho, and Francisco Serón. Pproc, an ontology for transparency in public procurement. *Se-*

-
- semantic Web*, 7(3):295–309, 2016.
- Alessandro Oltramari, Dhivya Piraviperumal, Florian Schaub, Shomir Wilson, Sushain Cherivirala, Thomas B Norton, N Cameron Russell, Peter Story, Joel Reidenberg, and Norman Sadeh. Privonto: A semantic framework for the analysis of privacy policies. *Semantic Web*, (Preprint):1–19, 2018.
- Monica Palmirani, Guido Governatori, Antonino Rotolo, Said Tabet, Harold Boley, and Adrian Paschke. Legalruleml: Xml-based rules and norms. In *Rule-Based Modeling and Computing on the Semantic Web*, pages 298–312. Springer, 2011.
- Monica Palmirani, Michele Martoni, Arianna Rossi, Cesare Bartolini, and Livio Robaldo. Pronto: Privacy ontology for legal reasoning. In *International Conference on Electronic Government and the Information Systems Perspective*, pages 139–152. Springer, 2018.
- Harshvardhan J Pandit, Kaniz Fatema, Declan OSullivan, and Dave Lewis. Gdprtext-gdpr as a linked data resource. In *European Semantic Web Conference*, pages 481–495. Springer, 2018.
- Livio Robaldo and Xin Sun. Reified input/output logic: Combining input/output logic and reification to represent norms coming from existing legislation. *Journal of Logic and Computation*, 27(8):2471–2503, 2017.
- Simon Steyskal and Axel Polleres. Defining expressive access policies for linked data using the odrl ontology 2.0. In *Proceedings of the 10th International Conference on Semantic Systems*, pages 20–23. ACM, 2014.
- Andre Valente, Jost Breuker, et al. A functional ontology of law. *Towards a global expert system in law*, pages 112–136, 1994.
- Robert van Kralingen. A conceptual frame-based ontology for the law. In *Proceedings of the First International Workshop on Legal Ontologies*, pages 6–17, 1997.