

Legal Ontologies and How to Choose Them: the InvestigatiOnt Tool*

Valentina Leone^{1,2}, Luigi Di Caro², and Serena Villata³

¹ University of Bologna, CIRSFD, Italy

² University of Turin, Italy

³ Université Côte d'Azur, CNRS, Inria, I3S, France

Abstract. Legal ontologies aim to provide a structured representation of legal concepts and their interconnections. These ontologies are then exploited to support information extraction and question answering in the legal domain. In addition, given the increasing importance of the Web of Data in public administration and in companies, being able to provide machine-readable legal information is becoming a valuable and desired contribution. The problem is that these ontologies are not always very transparent to end users, in particular if they lack legal knowledge. In this context, we present the InvestigatiOnt tool which aims to ease the interaction of end users with legal ontologies in order to spread the use of machine-processable legal information as well as its understanding.

1 Introduction

Ontologies are often at the basis of systems that support question answering, information extraction and knowledge modelling tasks. They are used to model the domain of knowledge for which a system is developed and the underlying concept structure. The design of ontology-based systems is usually assigned to computer scientists which need, in addition to the technical knowledge, a further knowledge about the domain for which the system is developed (e.g., economics, health care, law, agri-food sector). A particularly challenging domain is law, where concepts of increasing complexity are used and related to each other. In this context, there is the need to define tools able to support both developers and end-users towards a better understanding of the legal concepts expressed in the legal ontologies, so that an informed decision about the best ontology to select, depending on the target application, can be taken.

Based on these considerations, we propose a new tool called InvestigatiOnt. It aims to support developers and end-users with no expertise in law to choose the legal ontology that better suits the modelling requirements of a given legal domain of interest. In addition, InvestigatiOnt is not intended to only suggest an ontology to the user after she answers a list of questions, but to guide her in a learning process providing a better understanding about the distinctive features of each ontology. These features were identified through a detailed comparative analysis of the ontologies, inspired by [1,2].

Other tools make an effort to ease the use of legal ontologies to non-experts i.e., DALICC¹ and Licentia². They provide a support in the formulation of licences to

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

¹ www.dalicc.net

² www.licentia.inria.fr/

be applied on digital resources offering the possibility to specify the actions the licence should regulate. They also allow for compliance checking. InvestigatiOnt models a wider domain supporting the choice of an ontology in other legal fields than licenses.

2 The InvestigatiOnt tool

InvestigatiOnt is addressed to users which need to model a system concerning the legal domain. The legal field shows a lot of complexities for which InvestigatiOnt can provide a valuable decision support: (i) different legal systems (e.g., common law, civil law), (ii) different jurisdictions (e.g., local, national, international), (iii) complex interactions between norms (e.g., norms that express the obligation to be accomplished if another obligation is violated, or norms that express an exception to an obligation), (iv) the fixed structure of legal texts (e.g., their division in articles, paragraphs, definitions). InvestigatiOnt helps the user to deal with those challenging issues by offering her two types of service: the *visualization service* displays the information concerning an ontology as a diagram showing its dependencies with the other reused ontologies, while the *search service* suggests one or more ontologies suitable to meet the user requirements analyzing the answers she provided to a set of questions. The ontologies proposed to the user by this tool are those for which the ontology is available for download, and we concentrate on the more recently proposed ontologies (5-6 years ago at most). InvestigatiOnt deals with 12 ontologies which belong to six different legal fields:

1. *legal norms*: the ontologies model the norms as they could be found in the legal documents issued by local, national or international governments;
2. *policies*: the ontologies model the permitted, mandatory and prohibited actions that can be made on a digital or material asset;
3. *licences*: the ontologies model the actions allowed on a resource protected by the intellectual property right;
4. *legal documents representation/indexing*: the ontologies represent the text structure of legal documents and their topics;
5. *privacy in the GDPR*: the ontologies model the concepts involved in the new European General Data Protection Regulation (GDPR);
6. *tenders and public procurements*: the ontologies model the processes used by the public administration to find contractors to entrust with services or supplies.

The full list of the ontologies used in InvestigatiOnt is reported in Table 1. We now analyze the two services we developed to deal with these resources. A demo (recorded video) of InvestigatiOnt is available at <https://bit.ly/2Jid07I>.

The visualization service. It provides a visualization of the selected ontology as well as some information about the ontologies it reuses. More precisely, the user visualizes the following information for the selected ontology: (i) the *extended name of the ontology* and its acronym; (ii) the *year of publication* of the ontology for the first time; (iii) the *last update* of the resource (if this information is not displayed, the publication year coincides with the last update); (iv) the *licence* under which the ontology is made available for re-use, and (v) the *link to the ontology* in which the official description of the resource is made available, including the link for downloading it. In addition to these types of information, also a chart showing the dependencies among different ontologies is displayed. In particular, we considered two types of relations: *extends*,

Table 1: The list of ontologies used in InvestigatiOnt grouped by their legal field.

Field	Ontology	Main features
legal norms	LegalRuleML ³	modelling of the interpretations of a rule, tracking of the author of a document, modelling of the temporal evolution of norms
	NRV ⁴	extension of LegalRuleML, classification of deontic operators
policies	ODRL ⁵	modelling of roles of people involved in a policy, and of effects associated to the non compliance of different types of rules
	LDR ⁶	extends ODRL, models conditions of use of a LOD resource
licences	CC ⁷	modelling of the Creative Commons licence's concepts
	L4LOD ⁸	modelling of the actions which can be necessarily or possibly made and avoid on Linked Open Data
legal documents indexing	Eurovoc ⁹	semantic indexing of the documents issued by EU institutes
	ELI ontology ¹⁰	modelling of a set of metadata to allow the publication of legal documents of different EUs countries
Privacy in the GDPR	GDPRtEXT ¹¹	modelling of the concepts expressed in the GDPR as Linked Data, modelling of the structure of the GDPR text
	Bartolini et. al[3]	modelling of the GDPR and of the dependencies between rights and obligations with references between a right for the data producer and the corresponding obligation for the data user
tenders and procurements	LOTED2 ¹²	modelling of European tender notices as Linked Data
	PPROC ¹³	modelling of the public procurement process and the evolution of the contract, from its publication to its termination

indicating that the selected ontology recalls and extends the concepts of the other, and *re-uses*, indicating that the selected ontology reuses concepts from the other.

The search service. It aims to suggest to the user the ontology that better fits her requirements. To do this, the user is asked to answer a list of questions. Each question has a closed set of possible answers, and a response is required before moving to the next one. Each question aims to understand if and how a user needs to model a specific legal aspect of her domain of interest. InvestigatiOnt tries to understand the legal ontological commitment the user wants to assume. As we designed this system for users which are not familiar with the legal domain, questions are coupled with clarifying examples.

The first question asked to the user concerns the legal field that she needs to model for her purpose, i.e., the six legal fields listed in the introduction of this section. De-

³ <http://bit.ly/2sxpskV>

⁴ <http://ns.inria.fr/nrv/v1/nrvv1.html>

⁵ <http://www.w3.org/TR/odrl-model/>

⁶ <http://oeg-dev.dia.fi.upm.es/licensius/static/ldr/>

⁷ <http://creativecommons.org/ns>

⁸ http://ns.inria.fr/l4lod/v2/l4lod_v2.html

⁹ <http://eurovoc.europa.eu/drupal/?q=EuroVoc47>

¹⁰ <http://publications.europa.eu/mdr/eli/>

¹¹ <http://bit.ly/2xwjTZJ>

¹² <http://code.google.com/archive/p/loted2/source>

¹³ <http://contsem.unizar.es/def/sector-publico/pproc.html>

pending on the selected answer, the next questions that will be presented her will be different, in order to understand which ontology of the selected field is more suitable to fulfill the user requirements. These questions are the result of a careful study about similarities and differences among the ontologies of the same field performed together with legal experts. In particular, we explored two different ways to formulate a question: (i) the question recalls a feature belonging to one of the ontologies of the chosen field and the user is asked whether this feature is necessary for her modeling requirements (answer: *yes* or *no*), and (ii) the question asks the user to choose the way she wants the legal state-of-affairs to be modelled inside the ontology she is looking for (the kind of answer is more complex and it requires examples to ease the choice).

As the user goes on by answering the questions, the interface of *InvestigatiOnt* changes displaying further useful information. The first one is a track of the previous answers provided by the user. This is intended to help her to remember the selected answers, allowing her (if needed) to go back and change the answer to one or more questions. The second one is the information about the ontologies of the legal field chosen at the first step. In particular, every time a user answers a question, a score is assigned to one of the ontologies of the selected domain, depending on the selected answer. The basic idea is that each question tries to discriminate on a specific feature of the ontologies belonging to the selected field. A unitary score (graphically displayed as a star) is assigned to the ontologies and is added to the scores accumulated through the previous steps. Hence, through this mechanism, the user is supported in her selection process with a full transparency about the final recommendation provided by the system. As a result, the system shows a summary table where each row corresponds to a feature for which the user provided a response, while each column represents an ontology. A check symbol in a cell indicates that this ontology models that feature. This table provides an explainable recommendation to the user. This is particularly important when the scores between two or more ontologies are similar: the user has the possibility to evaluate the pros and cons of the choice of one resource rather than the other.

3 Conclusions

In this paper we presented *InvestigatiOnt*, a system prototype for an interactive and learning-oriented exploration of legal ontologies. Besides its concretization on the legal domain, the system may work on any domain and with additional visualization and exploration features. While most of the underlying motivations of the Semantic Web lie on the concept of *reuse*, the process of learning and (then) selecting existing ontological efforts is often strenuous due to domain complexity, subjective views, and specific application-oriented needs. Thus, the final aim of the presented system wants to be a stimulus for introducing and developing innovative tools to effectively enhance ontologies reuse and adaptation through interactive exploration and visualization features.

References

1. Casellas, N.: Legal ontology engineering: Methodologies, modelling trends, and the ontology of professional judicial knowledge. Volume 3. Springer Science & Business Media (2011)
2. Casanovas, P., Palmirani, M., Peroni, S., van Engers, T., Vitali, F.: Semantic web for the legal domain: the next step. *Semantic Web* 7(3) (2016)
3. Bartolini, C., Muthuri, R., Santos, C.: Using ontologies to model data protection requirements in workflows. In: *JSAI International Symposium on Artificial Intelligence*. (2015)