PrOnto: Privacy Ontology for Legal Reasoning

PrOnto:
Privacy Ontology for Legal Reasoning

Monica Palmirani[1]
Michele Martoni[2]
Arianna Rossi[3]
University of Bologna, CIRSFID, Via Galliera 3, 40121 Bologna, Italy

Cesare Bartolini[4]
Livio Robaldo[5]
SnT - Interdisciplinary Centre for Security, Reliability and Trust, Université du Luxembourg, JFK Building, 29, Avenue J.F. Kennedy, 1855 Luxembourg City, Luxembourg

This work was partially supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 690974 "MIREL: MIning and REasoning with Legal texts" and by the Luxembourg National Research Fund (FNR) CORE project C16/IS/11333956 "DAPRECO: DAta Protection REgulation COmpliance".

Abstract
The GDPR (GDPR, REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)) introduces the self-assessment of digital risks and the modulation of duties on the basis of the impact assessment analysis, including specific measures that intend to safeguard the data subject's human dignity and fundamental rights. Semantic web technologies and legal reasoning tools can support privacy-by-default and legal compliance. In this light, this paper presents a first draft of a legal ontology on the GDPR, called PrOnto, that has the goal of providing a legal knowledge modelling of the privacy agents, data types, types of processing operations, rights and obligations. The methodology used here is based on legal theory analysis joined with ontological patterns.

Keywords
Semantic web ; Legal reasoning ; Legal ontology ; Checking compliance

[1] monica.palmirani@unibo.it
[2] michele.martoni@unibo.it
[3] arianna.rossi15@unibo.it
[4] cesare.bartolini@uni.lu
[5] livio.robaldo@uni.lu

1 Introduction

The GDPR (*General Data Protection Regulation*) introduces a common legal framework for all the EU member states with the aim of harmonizing their privacy principles and the application of these principles inside the Digital Single Market. One of the main newly introduced instruments is the self-assessment of the digital risks and the modulation of the duties on the basis of the impact assessment analysis, including specific measures to safeguard the data subject's human dignity and fundamental rights. The audit and the compliance checking are instruments to guarantee *privacy-by-design* during software development (*ex-ante* phase) and the prompt detection of violations (*ex-post* phase) when they occur[6]. For this reason, semantic web and legal reasoning techniques can support the application of privacy-by-default principles in the day-byday operative tasks of public administrations, companies and non-profit organizations.

In this light, there is the urgent need to model a legal ontology of the privacy and data protection regulation, which must not be limited to the GDPR and which can be extended to other jurisdictions, in order to define the legal concepts in these legal frameworks and the relationships among them. This paper presents the first draft ontology on the GDPR, called PrOnto (Privacy Ontology), that aims to provide a legal knowledge modelling of the privacy agents, data types, processing operations, rights and obligations. The goal of this ontology is to support legal reasoning and check compliance by using defeasible logic theory (LegalRuleML standard [7] and SPINDle engine [8]), as opposed to exclusively improve information retrieval on the web.

[6]Casalicchio, E., Cardellini, V., Interino, G., Palmirani, M.: Research challenges in legal-ruleand QoS-aware cloud service brokerage Future Gener. Comput. Syst. 78, 211–223 (2016).

[7]Athan, T., Governatori, G., Palmirani, M., Paschke, A., Wyner, A.: LegalRuleML: design principles and foundations. In: Faber, W., Paschke, A. (eds.) Reasoning Web 2015. LNCS, vol. 9203, pp. 151–188. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-217680_6>

[8]Governatori, G., Hashmi, M., Lam, H.-P., Villata, S., Palmirani, M.: Semantic business process regulatory compliance checking using LegalRuleML. In: Blomqvist, E., Ciancarini, P., Poggi, F., Vitali, F. (eds.) EKAW 2016. LNCS (LNAI), vol. 10024, pp. 746–761. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-49004-5_48>

2 Related Work

Different authors from the semantic web community[9] have developed privacy ontologies for specific goals. For instance, the HL7 privacy ontology [10] is oriented to manage health data for electronic health records; others are oriented to secure messaging among automatic systems in the Internet of Things ecosystem, whereas others are oriented to manage the data flow in the linked open data environment or on the blockchain. However, there exists no legal ontology of privacy principles of the theory of law and foundational concepts that is able to support legal reasoning and check compliance. Those functionalities require a precise modelling of the rights and obligations using deontic operators and, at the same time, a modelling of the actors and the processing operations described in the normative prescriptions. For this reason, PrOnto takes inspiration from different existing ontologies and from the methodology of ontology design pattern[11]. We have used several other ontologies:

1. **ALLOT**: this ontology implements the Akoma Ntoso Top Level Classes (TLCs) as a formal OWL 2 DL and allows to connect the data and document classes with the FRBR ontology[12].
2. **FRBR**: FRBR is an ontology that implements the FRBR model[13].

---

[9]Ashley, K.: Artificial Intelligence and Legal Analytics New Tools for Law Practice in the Digital Age. Cambridge University Press, Cambridge (2017) ; Gharib, M., Giorgini, P., Mylopoulos, J.: Towards an ontology for privacy requirements via a systematic literature review. In: Mayr, H.C., Guizzardi, G., Ma, H., Pastor, O. (eds.) ER 2017. LNCS, vol. 10650, pp. 193–208. Springer, Cham (2017). <https://doi.org/10.1007/9783-319-69904-2_16> ; <http://www.w3.org/Privacy/> ; Proceedings of the 5th Workshop on Society, Privacy and the Semantic Web - Policy and Technology (PrivOn2017) co-located with 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria, 22 October 2017. CEUR Workshop Proceedings 1951, CEUR-WS.org 2017. <http://events.linkeddata.org/ldow2011/papers/ldow2011-paper01sacco.pdf> ; Samavi, R., Consens, M.P.: Publishing privacy logs to facilitate transparency and accountability. J. Semant. Web 50, 1–20 (2018)

[10]<http://wiki.hl7.org/index.php?title=Security_and_Privacy_Ontology> ; <http://www.hl7.org/implement/standards/product_brief.cfm?product_id=348>

[11]Gandon, F., Governatori, G., Villata, S.: Normative requirements as linked data. In: JURIX2017. IOS Press (2017)

[12]Barabucci, G., Cervone, L., Di Iorio, A., Palmirani, M., Peroni, S., Vitali, F.: Managing semantics in XML vocabularies: an experience in the legal and legislative domain. In: Proceedings of Balisage 2009 (2010).

[13]IFLA Study Group on the FRBR: Functional requirements for bibliographic records (2009).

3. **LKIF Core**: Action.owl is an ontology that represents actions in general, i.e., processes that are performed by an agent. We use in particular lkif:Agent to model lkif:Organization and lkif:Person.

4. **LKIF Core**: Role.owl is an ontology to describe typologies of roles (epistemic roles, functions, person roles, organisation roles). We use in particular lkif:Role[14].

5. The **Publishing Workflow Ontology** (PWO) is a simple ontology written in OWL 2 DL for the characterization of the main stages in the workflow associated with the publication of a document (e.g., being written, under review, XML capture, page design, publication on the Web). We reuse the workflow pattern to model the different types of processing of personal data[15].

6. **Time-indexed Value in Context** (TVC) is an ontology pattern that allows to describe scenarios in which someone (e.g., a person) has a value (e.g., a particular role) during a particular time and for a particular context. We use this portion of ontology to connect the event with value, context and time parameters[16].

7. **Time Interval** (TI) is an ontology design pattern that enables the description of periods of time that are characterised by a starting date and an ending date. We use this ontology to manage the time interval[17].

3 Methodology: MeLOn

We developed PrOnto by using an interdisciplinary approach called MeLOn (Methodology for building Legal Ontology), which has been already used with success to develop several legal ontologies. The MeLOn methodology was built to design legal ontologies, considering the great difficulties that legal experts encounter when they must define a model of the reality using the ontological techniques. Protégé was used frequently in the past in the legal community, but with the result to produce a large number of classes, one for each legal term, because the legal expert is not usually familiar with the modelization of the reality

[14]Breuker, J.A.P.J., et al.: OWL ontology of basic legal concepts (LKIF-Core). Estrella: Deliverable 1.4., AMSTERDAM, UVA, 2007, p. 138 (2007)

[15]Gangemi, A., Peroni, S., Shotton, D., Vitali, F.: The publishing workflow ontology (PWO), ISO (2016). <http://www.semantic-web-journal.net/content/publishing-workflow-ontologypwo>.

[16]Peroni, S., Palmirani, M., Vitali, F.: UNDO: the United Nations system document ontology.In: d'Amato, C., et al. (eds.) ISWC 2017 Part II. LNCS, vol. 10588, pp. 175–183. Springer, Cham (2017). <https://doi.org/10.1007/978-3-319-68204-4_18>.

[17]*Ivi.*

using classes, relationships and attributes. The Glossary method is too language-oriented. The foundational approach is too abstract and too little applicative, even if DOLCE[18] is used as skeleton for the final checking.

The MeLOn methodology is composed of ten steps that can be recursively applied:

1 Describe the goal of the ontology. In this step, the team describes the research questions that the ontology intends to cope with. It is also important to select two or three use-cases where the ontology is helpful. For PrOnto we defined the following goals:

(i) to model data protection legal norms starting from legal texts but including also social norms, practitioner opinions or social behaviours;

(ii) to build a legal ontology that is usable for legal reasoning;

(iii) to build a legal ontology that is usable for web of data and information retrieval.

2 Evaluation indicators. We define some parameters/indicators to evaluate the ontology according to the goals (step 1). In the PrOnto ontology, we selected the following criteria based on the existing state of the art[19];

(i) **coherence**: the axioms of the ontology can't create inconsistency or contradictions;

(ii) **completeness**: the domain is adequately covered by the ontology and the main concepts are included;

(iii) **efficiency**: the ontology is technically sound, concise and the reasoning is computable in reasonable time, and it is based on patterns;

(iv) **effectiveness**: the ontology covers the most important queries about the domain and the end users find it helpful to resolve applicative situations;

(v) **usability**: the end users find the ontology clear, understandable, easy to use, close to the main terminology used inside of the community, sefl-explained.

---

[18]Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L.: Sweetening ontologies with DOLCE. In: Gómez-Pérez, A., Benjamins, V.R. (eds.) EKAW 2002. LNCS (LNAI), vol. 2473, pp. 166–181. Springer, Heidelberg (2002). <https://doi.org/10.1007/3-540-458107_18>.

[19]Bandeira, J., Bittencourt, I., Espinheira, P., Isotani, S.: FOCA: a methodology for ontology evaluation, arxiv (2016).

(vi) **agreement**: the grade of agreement and acceptance of the ontology in the legal expert community.

3 State of the art survey. We have checked the state of the art in order to reuse existing ontologies, ontology patterns[20], and other existing domain vocabularies.

4 List all the relevant terminology. We produce a glossary with the most relevant legal terms extracted from normative documents, case-law, contracts, or any other legal source. In particular, we included all the legal definitions.

5 Use usable tools. We use tools that are close to the legal experts such as tables or UML diagrams in order to model the knowledge-base of the legal domain. Legal experts can use the Graffoo tool[21] that allows to use graphical instruments and to transform the UML into OWL/XML serialization.

6 Refine and optimize. The serialization into OWL by Graffoo[22] [8, 21] or UML is not optimal for the efficiency and the coherence, therefore the axioms are added manually by an ontology expert in order to check the coherence.

7 Test the output. The ontology is tested by legal experts using a web interface in order to evaluate the completeness, effectiveness and usability.

8 Evaluate the ontology. We use the OntoClean method to polish the ontology and apply the criteria of point 2 to provide metrics. A set of SPARQL queries are prepared and the output is measured.

---

[20]Hitzler, P., Gangemi, A., Janowicz, K., Krisnadhi, A.A., Presutti, V.: Ontology Engineeringwith Ontology Design Patterns: Foundations and Applications. IOS Press, Amsterdam (2016).

[21]<http://www.essepuntato.it/graffoo/> ;
<http://www.yworks.com/en/products/yfiles/yed>.

[22]Falco, R., Gangemi, A., Peroni, S., Shotton, D., Vitali, F.: Modelling OWL ontologies with Graffoo. In: Presutti, V., Blomqvist, E., Troncy, R., Sack, H., Papadakis, I., Tordai, A. (eds.) ESWC 2014. LNCS, vol. 8798, pp. 320–325. Springer, Cham (2014). <https://doi.org/10. 1007/978-3-319-11955-7_42> ; Peroni, S.: A simplified agile methodology for ontology development. In: Dragoni, M., Poveda-Villalón, M., Jimenez-Ruiz, E. (eds.) OWLED/ORE -2016. LNCS, vol. 10161, pp. 55–69. Springer, Cham (2017). <https://doi.org/10.1007/978-3-319-54627-8_5>.

9 Publish the document with the LODE tool[23].

10 Collect feedbacks from the community in order to reach the agreement criteria.

The method must be repeated at least three times and transparently published online.

4 PrOnto Modules

PrOnto consists of different modules: (i) documents and data, (ii) actors and roles, (iii) processing and workflow, (iv) legal rules and deontic formula, (v) purposes and legal bases (**Fig. 1**).

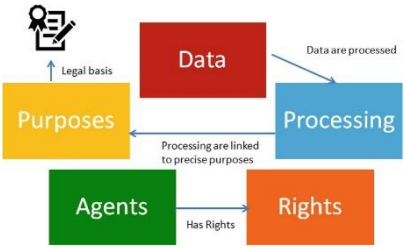

Fig. 1. Modules of PrOnto ontology

Some document and data are referred to the data subject. Data subject is a *role* of an *agent* (physical person). Data is processed following a given *workflow* plan of actions. When executed, each action assumes specific temporal parameters (e.g., the processing's interval of time), context (e.g., jurisdiction where the data processing is carried out), and value (e.g., place where the data processing is performed). The data processing must be performed according to a *legal basis* that provides the lawfulness of the processing. Each *processing* activity involves a controller, a processor, and other actors. Each actor has obligations or rights, for instance the data subject has rights related to the data protection. These rights and obligations are linked to documents where the norms appear: terms of use, information, privacy policies, consent forms.

[23]<http://www.essepuntato.it/lode> ; Peroni, S., Shotton, D., Vitali, F.: The live OWL documentation environment: a tool for the automatic generation of ontology documentation. In: ten Teije, A., et al. (eds.) EKAW 2012. LNCS (LNAI), vol. 7603, pp. 398–412. Springer, Heidelberg (2012). <https://doi.org/10. 1007/978-3-642-33876-2_35>.

4.1 Data and Document

Data protection involves data and documents in a twofold manner: data are the object of the regulation and the target of its protection, and also the source of information to regulate the relationships between the different agents (e.g., controller, processor, etc.) using privacy, informed consent, contracts, codes of conduct, law, case-law and any other legal document. The data and the documents are documental sources; using the FRBR ontology, we model their representations over time by reusing a robust design pattern already adopted for the publication process[24]. Data are defined in categories according to the GDPR: personal data, non-personal data, anonymized data, pseudonymised data (**Fig. 2**).

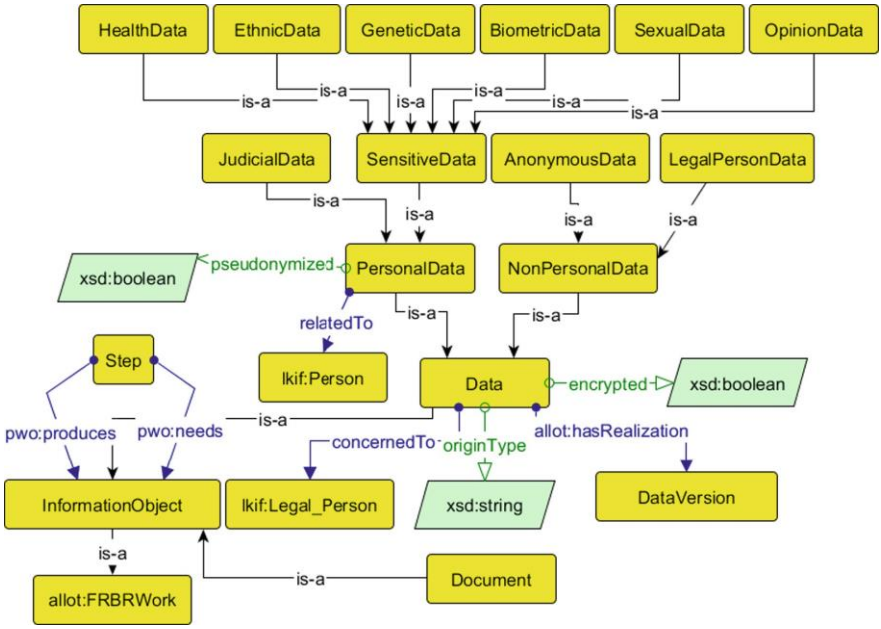

Fig. 2. Document and data module

4.2 Agent and Role

One of the most frequent errors in legal ontology design is to confuse agents and roles. In PrOnto we clearly distinguish the two classes. Physical persons and organizations are agents, but we include into the agent class also IT organizations or artificial intelligence and software or robots. An agent could play multiple roles related to different processing activities and contexts. Additionally, a controller could act as processor

[24]Peroni, S., Shotton, D.: The SPAR ontologies. To appear in Proceedings of the 17th International Semantic Web Conference, ISWC2108 (2018 under publication). <https://w3id. org/spar/article/spar-iswc2018/>

or third party with respect to a separate processing. Each role is fixed in a given period of time that is joined with the time version of the dataset and the duration of the data processing. The role is authorized by an event that assigns it to the agent (see **Fig. 3**). The role is modelled in subclasses like DPO (data protection officer), controller, processor, third party, representative, recipient, data subject, supervisory authority, Member State. Other roles are defined by the deontic legal rules such as bearer or counter party.

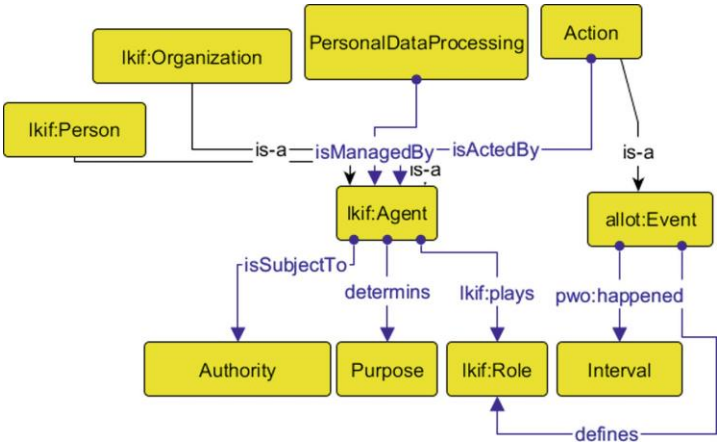

Fig. 3. Agent and role module

4.3 Data Processing

When we model human activities, we need to model workflows as a sequence of steps that uses some resources in input and produces some outcomes. However, a workflow is composed of two parts: the plan to do something (e.g., workflow) and the concrete sequence of actions actually performed (e.g., execution of the workflow). In the GDPR, it is especially important to distinguish the plan (e.g., Impact Assessment Plan made of steps) from the real execution (e.g., data breach event and counter measurement enacted), which is constituted by a set of actions. Especially in the compliance checking scenario, there is the need to have a plan that conforms to the law and to provide counter measures in case of violation during the actual execution (e.g., remedies). For this reason, we have used the Publishing Workflow Ontology (PWO) as the basis to model the data processing ontology module. PWO incudes workflow and executed workflow. *PersonalDataProcessing* is a subclass of Workflow with several attributes: *transparency*, *fairness*, *lawfulness* that are Boolean value that a legal reasoning process could set up. Personal data processing is also planned for being eligible for a given period of time (*isValid*), also in accordance with the purpose (*isBasedOn*).

PersonalDataProcessingExecution is a subclass of *WorkflowExecution*. The workflow execution involves actions. The actions[25] are a kind of *event* that are described by temporal parameters (e.g., interval) and context values (Time-indexed Value in Context - TVC). The Action class in PrOnto also has an important attribute for storing the status of *breachness*: the action is prone to configure a data breach event. One of the values of the action is the place where the event occurs (e.g., within the EU borders) and the *jurisdiction* (e.g., Regional competence). Other values and statuses can be added in order to enrich the context description (see **Fig. 4**).

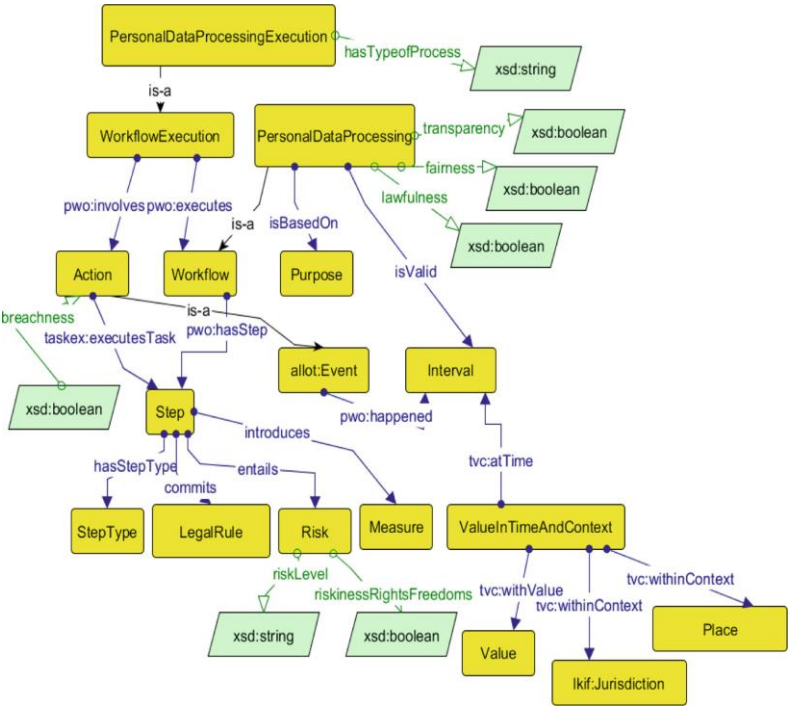

Fig. 4. Workflow and processing module

For instance, we take the category of all the actions that produce a "deletion" according with the Article 17 of GDPR. Technically speaking, it is not easy to isolate the exact moment and level of deletion (e.g., logical deletion or physical erasure – see **Fig. 5**), but under the legal point of view we can include in this category the following behaviours: a temporary deletion, a permanent deletion including the backup copies in cloud computing, destruction of the physical device, anonymisation of

[25]Abrams, M.: The origins of personal data and its implications for governance. <https://papers. ssrn.com/sol3/papers.cfm?abstract_id=2510927>.

the data, and finally the pseudoanonymisation of data with double password access and kept in a secure place (e.g., safe). However, there are situations in which it is difficult to ensure a total erasure (e.g., blockchain), and the anonymisation techniques do not guarantee 100% security of de-identification[26]. For these reasons, PrOnto distinguishes between different levels of delete actions: PermanentErasure, Destroy e Anonymise. The deletion action is also activated when the processing expires. When the purpose and the valid period expires, the ontology can execute the deletion action.

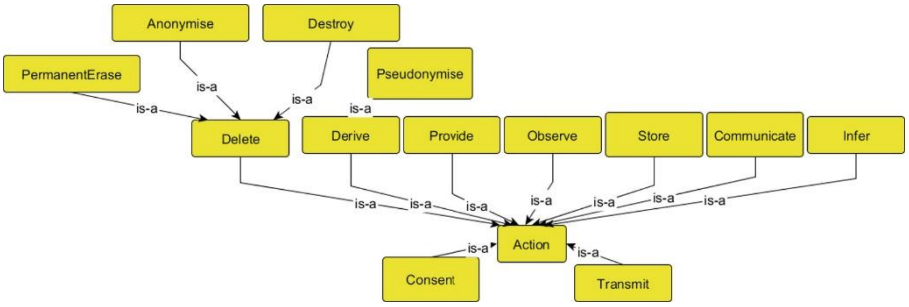

Fig. 5. Action module

4.4 Purposes and Legal Basis

The GDPR permits the processing of personal data only in the light of several lawful purposes. The purposes must be supported by a legal basis (Article 6 – Lawfulness of processing). For this reason, we have introduced a *lawfulness* status as a Boolean data property of the *PersonalDataProcessing* class. Each personal data processing is based on a *Purpose*. In this way, a rule engine, based for instance on a rule language like LegalRuleML, can return this value after the rule reasoning process (**Figs. 6** and **7**).

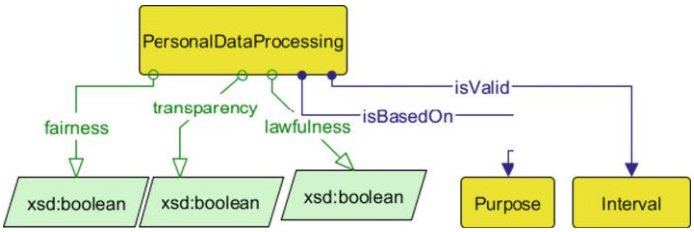

Fig. 6. Lawfulness status and legal basis relationship

[26] Deleting personal data. <https://ico.org.uk/media/for-organisations/documents/1475/deleting_ personal_data.pdf>.

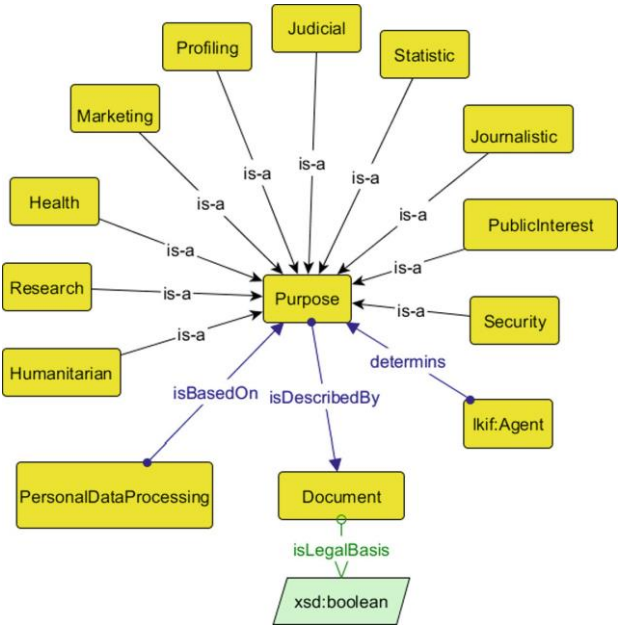

Fig. 7. Purpose class and subclasses

4.5 Deontic Operators

The modelling of legal norms needs deontic operators such as right, obligation, permission and prohibition. From the point of view of the GDPR, it is very relevant to also include violation/compliance as the status where an obligation or a prohibition is violated or is compliant. The deontic operators are connected to temporal parameters, and to a jurisdiction as well, in case some rights are effective only in a certain domestic regulation. This part of the PrOnto ontology allows us to model the necessary predicates to implement legal rules. This module is an extension of the LegalRuleML meta model, which allows us to synchronize the legal rule language modelling with the ontology.

Each step commits a LegalRule that is made up of Deontic Specifications (**Fig. 8**). The Right and Obligation classes are detailed in subclasses according to the GDPR. Right is connected to a permission. In this manner, we can track the permission connected with a specific right such as the right to access (e.g., permission to use a PET – Privacy-enhancing technology), whilst obligation is connected to violation or compliance. We are thus able to make queries like the following: give me all the obligations of the controller (X) that were violated in a given interval $[t_x, t_y]$ (**Figs. 9** and **10**).

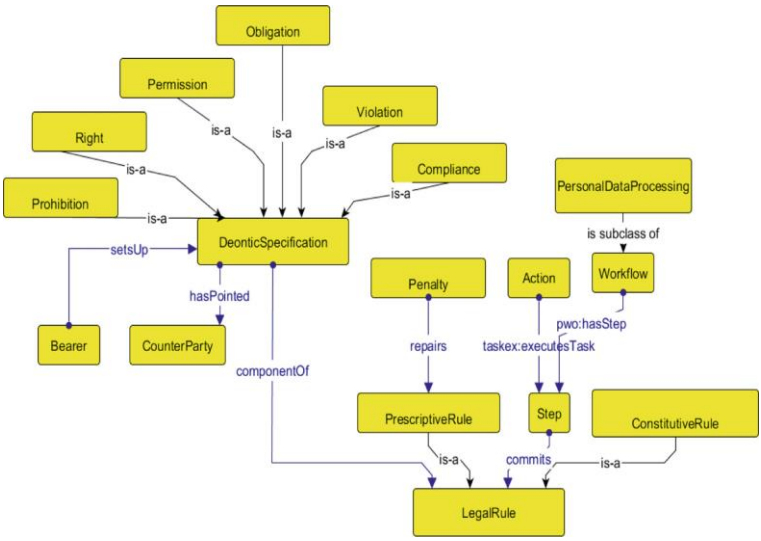

Fig. 8. LegalRule module

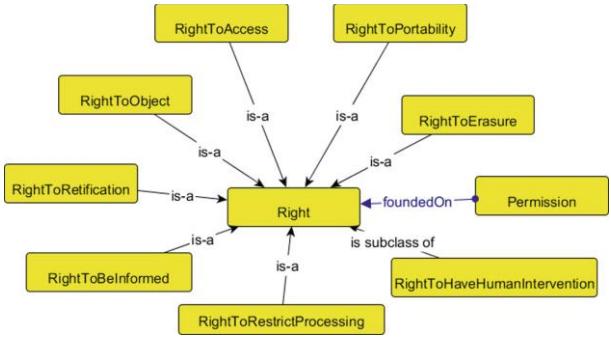

Fig. 9. Right classes

The ontology in this module also intends to model the relationships between deontic rules, actors' rights and obligations, obligations and permissions, and violation/compliance. This modelling allows to populate the ontology, or to create RDF triples, in order to perform queries like the following: "give me all the data processing that has been violated by some actors in a given time". This knowledge is processed by the rule engine, but transformed into individuals in the ontology (or RDF triples) without the need to query to the rule engine each time.

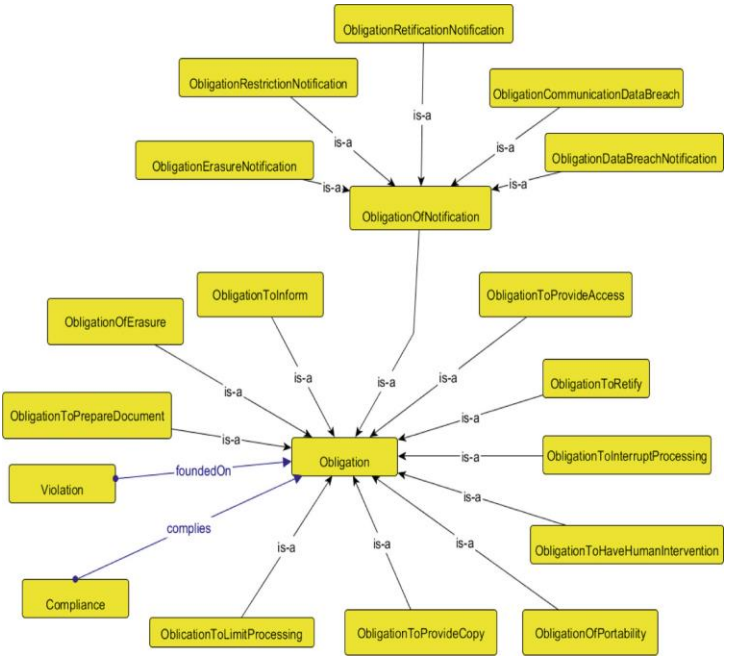

Fig. 10. Obligation classes

5 Evaluation

The evaluation is carried out inside the Cloud4EU European project PCP[27] that intends to provide legal compliance checking systems for eGovernment services that are delivered across the cloud. We are currently in the phase of testing PrOnto on three different scenarios related to school services. PrOnto is also used inside the MIREL European project[28] and the DAPRECO Luxembourgish project[29].

An example of the use of PrOnto is presented hereafter.

[27]<http://www.agid.gov.it/cloudforeurope>.
[28]<http://www.mirelproject.eu/>.
[29] <https://wwwen.uni.lu/research/fstc/computer_science_and_communications_research_unit/research_ projects/data_protection_regulation_compliance>.

a. Give me all the personal data processing performed by company X in the role of controller valid in [t_1, t_2].	SELECT ?pdp
	WHERE {
	?pdp :isManagedBy _:c .
	[lkif:plays _:c ;
	rdfs:label "X"] .
	?pdp :isValid [
	time:hasBeginning [rdfs:label "t1"] ;
	time:hasEnd [rdfs:label "t2"]
	] .
	}

b. Give me all the communications connected with of a given step K in the PersonalDataProcessing.	SELECT ?a ?pdp
	WHERE {
	?a a :Action .
	?a taskex:executesTask _:s .
	?pdp pwo:hasStep _:s .
	_:s rdfs:label "K" .
	}

The previous queries produce important results to check the GDPR obligations and facilitating a dynamic self-assessment. We suppose that a software manages documentation, registry of processing, DPIA information, etc. (e.g., software provided by the French CNIL – Commission Nationale de l'Informatique et des Libertés[30]). If such a software is connected with PrOnto ontology, we can check for GDPR compliance throughout all the lifecycle of the personal data, using advanced legal reasoning tools or SPARQL end-points.

6 Conclusions and Future Work

Several privacy ontologies exist (e.g., HL7 for eHealth, PPO for Linked Open Data, OdrL for modelling rights, etc.) in the state of the art but are not integrated with deontic logic models usable for legal reasoning. PRONTO intends to integrate different levels of semantic representation: document and data modelling to support the semantic web information retrieval, in particular Linked Open Data (e.g., SPARQL queries); workflow and processing to support the planning of privacy policy and possibly also BPMN modelling for system design (e.g., privacy-by-design); rights and obligations to enable the legal reasoning using rule languages (e.g., LegalRuleML and compliance checking); human-centric approaches to favour the visualization and the

[30] <https://www.cnil.fr/en/open-source-pia-software-helps-carry-out-data-protection-impact-assesment>.

presentation of the privacy-related legal principles and concepts in different contexts and towards different targets.

This is a long-term research. We intend to proceed with the modelling and optimization of the formal ontology and to evaluate it with a large number of use-cases. In the meantime, we believe that such an ontology has to be negotiated with a large community, in order to create a consensus and to place those results into a standardization body for the future governance (e.g., OASIS, W3C). In the future, it is also necessary to develop specific profiles, one for each specific national law, or by thematic domain (e.g., Privacy in IoT, Privacy in AI, etc.).

References

Abrams, M.: The origins of personal data and its implications for governance. <https://papers. ssrn.com/sol3/papers.cfm?abstract_id=2510927>

Ashley, K.: Artificial Intelligence and Legal Analytics New Tools for Law Practice in the Digital Age. Cambridge University Press, Cambridge (2017)

Athan, T., Governatori, G., Palmirani, M., Paschke, A., Wyner, A.: LegalRuleML: design principles and foundations. In: Faber, W., Paschke, A. (eds.) Reasoning Web 2015. LNCS, vol. 9203, pp. 151–188. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-217680_6>

Bandeira, J., Bittencourt, I., Espinheira, P., Isotani, S.: FOCA: a methodology for ontology evaluation, arxiv (2016)

Barabucci, G., Cervone, L., Di Iorio, A., Palmirani, M., Peroni, S., Vitali, F.: Managing semantics in XML vocabularies: an experience in the legal and legislative domain. In: Proceedings of Balisage 2009 (2010)

Breuker, J.A.P.J., et al.: OWL ontology of basic legal concepts (LKIF-Core). Estrella: Deliverable 1.4., AMSTERDAM, UVA, 2007, p. 138 (2007)

Casalicchio, E., Cardellini, V., Interino, G., Palmirani, M.: Research challenges in legal-ruleand QoS-aware cloud service brokerage Future Gener. Comput. Syst. 78, 211–223 (2016)

Falco, R., Gangemi, A., Peroni, S., Shotton, D., Vitali, F.: Modelling OWL ontologies with Graffoo. In: Presutti, V., Blomqvist, E., Troncy, R., Sack, H., Papadakis, I., Tordai, A. (eds.) ESWC 2014. LNCS, vol. 8798, pp. 320–325. Springer, Cham (2014). <https://doi.org/10. 1007/978-3-319-11955-7_42>

Gandon, F., Governatori, G., Villata, S.: Normative requirements as linked data. In: JURIX2017. IOS Press (2017)

Gangemi, A., Peroni, S., Shotton, D., Vitali, F.: The publishing workflow ontology (PWO), ISO (2016). <http://www.semantic-web-journal.net/content/publishing-workflow-ontologypwo>

Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L.: Sweetening ontologies with DOLCE. In: Gómez-Pérez, A., Benjamins, V.R. (eds.) EKAW 2002. LNCS (LNAI), vol. 2473, pp. 166–181. Springer, Heidelberg (2002). <https://doi.org/10.1007/3-540-458107_18>

Gharib, M., Giorgini, P., Mylopoulos, J.: Towards an ontology for privacy requirements via a systematic literature review. In: Mayr, H.C., Guizzardi, G., Ma, H., Pastor, O. (eds.) ER 2017. LNCS, vol. 10650, pp. 193–208. Springer, Cham (2017). <https://doi.org/10.1007/9783-319-69904-2_16>

Governatori, G., Hashmi, M., Lam, H.-P., Villata, S., Palmirani, M.: Semantic business process regulatory compliance checking using LegalRuleML. In: Blomqvist, E., Ciancarini, P., Poggi, F., Vitali, F. (eds.) EKAW 2016. LNCS (LNAI), vol. 10024, pp. 746–761. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-49004-5_48>

Hitzler, P., Gangemi, A., Janowicz, K., Krisnadhi, A.A., Presutti, V.: Ontology Engineeringwith Ontology Design Patterns: Foundations and Applications. IOS Press, Amsterdam (2016)

<http://wiki.hl7.org/index.php?title=Security_and_Privacy_Ontology>

<http://www.hl7.org/implement/standards/product_brief.cfm?product_id=348>

<http://www.w3.org/Privacy/>

IFLA Study Group on the FRBR: Functional requirements for bibliographic records (2009).

<http://www.ifla.org/publications/functional-requirements-for-bibliographic-records. Accessed 7 May 2017>

Peroni, S., Palmirani, M., Vitali, F.: UNDO: the United Nations system document ontology.In: d'Amato, C., et al. (eds.) ISWC 2017 Part II. LNCS, vol. 10588, pp. 175–183. Springer, Cham (2017). <https://doi.org/10.1007/978-3-319-68204-4_18>

Peroni, S., Shotton, D., Vitali, F.: The live OWL documentation environment: a tool for theautomatic generation of ontology documentation. In: ten Teije, A., et al. (eds.) EKAW 2012. LNCS (LNAI), vol. 7603, pp. 398–412. Springer, Heidelberg (2012). <https://doi.org/10. 1007/978-3-642-33876-2_35>

Peroni, S.: A simplified agile methodology for ontology development. In: Dragoni, M., Poveda-Villalón, M., Jimenez-Ruiz, E. (eds.) OWLED/ORE -2016. LNCS, vol. 10161, pp. 55–69. Springer, Cham (2017). <https://doi.org/10.1007/978-3-319-54627-8_5>

Peroni, S., Shotton, D.: The SPAR ontologies. To appear in Proceedings of the 17thInternational Semantic Web Conference, ISWC2108 (2018 under publication). <https://w3id. org/spar/article/spar-iswc2018/>

Proceedings of the 5th Workshop on Society, Privacy and the Semantic Web - Policy andTechnology (PrivOn2017) co-located with 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria, 22 October 2017. CEUR Workshop Proceedings 1951, CEUR-WS.org 2017. <http://events.linkeddata.org/ldow2011/papers/ldow2011-paper01sacco.pdf>

Samavi, R., Consens, M.P.: Publishing privacy logs to facilitate transparency and accountability. J. Semant. Web 50, 1–20 (2018)

Deleting personal data. https://ico.org.uk/media/for-organisations/documents/1475/deleting_ personal_data.pdf