



An Agile Approach to Validate a Formal Representation of the GDPR

Cesare Bartolini^{1(✉)}, Gabriele Lenzini¹, and Cristiana Santos^{2,3}

¹ Interdisciplinary Centre for Security, Reliability and Trust (SnT),
University of Luxembourg, Luxembourg, Luxembourg
{cesare.bartolini,gabriele.lenzini}@uni.lu

² Research Centre for Justice and Governance (JusGov), School of Law,
University of Minho, Braga, Portugal

³ Université Toulouse Capitole, Toulouse, France
cristiana.teixeirasantos@gmail.com

Abstract. Modelling in a knowledge base of logic formulæ the articles of the GDPR enables a semi-automatic *reasoning* of the Regulation. To be legally substantiated, it requires that the formulæ express validly the legal meaning of the Regulation’s articles. But legal experts are usually not familiar with logic, and this calls for an interdisciplinary validation methodology that bridges the communication gap between formal modelers and legal evaluators. We devise such a validation methodology and exemplify it over a knowledge base of articles of the GDPR translated into Reified I/O (RIO) logic and encoded in LegalRuleML. A pivotal element of the methodology is a human-readable intermediate representation of the logic formulæ that preserves the formulæ’s meaning, while rendering it in a readable way to non-experts. After being applied over a use case, we prove that it is possible to retrieve feedback from legal experts about the formal representation of Art. 5.1a and Art. 7.1. What emerges is an agile process to build logic knowledge bases of legal texts, and to support their public trust, which we intend to use for a logic model of the GDPR, called DAPRECO knowledge base.

Keywords: General Data Protection Regulation · Data protection · Compliance · Legal validation · Usability

1 Introduction

Artificial Intelligence (AI) applied to the legal domain can serve a number of purposes to improve the efficiency of legal services and the predictability of the application of the law. Some illustrative examples of legal AI applications are evinced by existing digital services such as: search engine for retrieving legal

Bartolini and Lenzini are supported by the FNR CORE project C16/IS/11333956 “DAPRECO: DAta Protection REgulation COmpliance”.

sources; online dispute resolution; assistance in drafting needs; predictive analysis; categorization of contracts and detection of incompatible contractual clauses; “chatbots” to support litigants; and legal reasoning and decision-making. Moreover, AI compliance tools can help to identify the laws and regulations a certain business activity is subject to, assisting undertakings in establishing legally-compliant processes, and easing the verification of compliance by auditors and enforcement bodies.

Instantiations of AI-enabled tools for legal compliance within data protection seems particularly pertinent. Notably, the new legal landscape reshaped by the General Data Protection Regulation (GDPR), coupled with the heavy fines that supervisory authorities are entitled to issue in case of data breach, calls for a need to ensure compliance for data processing activities. Herein, data controllers could use AI compliance tools designed to help them assuring accountability and compliant management processes, whilst diminishing the risks of violating provisions and incurring into fines.

A critical facet of such automation is the need to build executable rules for a computer-assisted compliance system. In previous research, the authors have proposed a complete model of the GDPR for legal reasoning and legal compliance [2, 15–17]. This model comprises three components: (i) the legal text in Akoma Ntoso format; (ii) an ontology of legal concepts concerning privacy and data protection; and (iii) a knowledge base of data protection rules. This last component, called the Data Protection Regulation Compliance (DAPRECO) Knowledge Base¹, currently under development, contains the General Data Protection Regulation (GDPR) provisions modeled in Reified Input/Output (RIO) logic [20]. Built to contain natural language interpretations of these provisions, due to the logic’s defeasible nature, the DAPRECO Knowledge Base can be updated with successive and more authoritative legal interpretations [2]. Accordingly, the Knowledge Base needs to be adequately validated before it can perform in a real-world environment. Such a pragmatical stand is demanded, since “for developers, as contrasted to researchers, the issue is not whether the resulting rule base is complete or even accurate or self-modifying – but whether the rule base is sufficiently accurate to be useful” [5] when it is moved out from the research laboratory and into the marketplace [23]. However, as is widely acknowledged in literature [11, 19, 22], testing legal Artificial Intelligence (AI) systems is a difficult task because approaches reveal coder-dependency, and it is complex to emulate the “art-of-the-experts” [6]. With ongoing maturity in the field of AI and Law, the need for an easily accessible *interdisciplinary validation methodology* comes into play [10].

Legal Validation. The concept of validation refers to the determination of the correctness of the system with respect to user needs and requirements [22]. Legal validation is “needed to verify the correctness of the output of the system in relation to the knowledge of the legal domain it covers”, “the guarantee of the one-to-one relation between analysis and representation” [12]. Such a method

¹ The name DAPRECO comes from DAta PRotection REgulation COmpliance, the name of the CORE-FNR project that supported this research.

would assist legal professionals framing an evaluation of an AI-legal system and help IT experts understand the validation requirements of legal professionals [11]. As “algorithmic representations of law are typically very poor as regards their transparency”, “one cannot begin to devise an algorithm to apply legal provisions without determining first its intended purpose and by whom it will be used” [21]. Thus, validating a legal model requires that the formalization used is *understandable and accessible*. Consequently, the methodology should be driven by *usability* considerations in the adopted criteria, and validation tests (through user acceptance surveys or questionnaires) [22], as foreseen in the current work.

This validation quadrant also holds for our domain modeling of the GDPR. We believe that this endeavour to formalize articles in a logic formalism requires a methodology supporting its legal validation. The validation phase should not be postponed till the moment when the whole GDPR is formalized, for as detecting possible unsound conclusions at such late juncture would amount to a very expensive step, likely inspiring distrust in the whole framework. A more *agile* process is advisable and was adopted to validate the legal soundness of any formula from the moment in which they were added to the Knowledge Base, thus assisting incrementally and concomitantly the modeler.

Contribution. This paper builds on two workshop articles [3, 4]. The contribution is a methodology aiming to capture *informed* feedback on the legal validity of the DAPRECO Knowledge Base’s representing the meaning(s) of the articles of the GDPR. A decisive element of the methodology is a human-readable break-down of a RIO logic formula. Once the customizable human-readable representation has been assessed as understandable, increasing our confidence on it to be an eligible candidate to validate the formalized GDPR articles, we proceed further and show that the methodology is effective in gathering feedback of legal experts on the legal validity of the representation of the GDPR articles, so as to provide quality assurance of our methodology as a whole.

This paper reports fully on the study, comments on the methodology, and on the usability experiments, pointing out the limitations and future work.

2 Related Work

Some discussion within the AI and Law community [10, 11, 22] – specifically amidst the Proceedings of the International Conference on Artificial Intelligence and Law works (ICAIL), and later through the Journal of Artificial Intelligence and Law contributions (JAIL), – concerned qualitative evaluation methodologies suitable for legal domain systems, and the best practices through which AI and Law researchers could frame the assessment of the performance of their works, both empirical and theoretical. For example, performance evaluation is emphasized and compared to known baselines and parameters, using publicly available datasets whenever possible [8, 9].

A set of six categories was compiled to define the broad types of evaluation found therefrom. They include the following assessments: i. *Gold Data*: evaluation performed with respect to domain expert judgments (e.g., classification

measurements or measures on accuracy, precision, recall, F-score, etc.); ii. *Statistical*: evaluation performed with respect to comparison functions (e.g., unsupervised learning; cluster internal-similarity, cosine similarity, etc.); iii. *Manual Assessment*: performance is measured by humans via inspection, assessment, review of output; iv. *Algorithmic*: assessment made in terms of performance of a system, such as a multi-agent system; v. *Operational-Usability*: assessment of a system's operational characteristics or usability aspects; vi. *Other*: those systems with distinct forms of evaluation not covered in the categories above (task-based, conversion-based, etc.). In *our case*, we combined the following types of evaluation: gold data (i.), manual assessment (iii.) and operational-usability (v.).

Some authors [10] developed the Context Criteria Contingency-guidelines Framework (CCCF) for evaluating Legal Knowledge Based System (LKBS). Within this framework, the quadrant criteria pertinent to the purposes of this paper are herewith mentioned. The *User Credibility* quadrant refers to credibility and acceptability of a system at the individual level. It comprises three main branches associated with *user satisfaction*, *utility* (usefulness or fitness for purpose) and *usability* (ease of use). The usability branch is further decomposed into branches associated with operability, understandability, learnability, accessibility, flexibility in use, and with other human factors and human computer interface issues. The *Verification and Validation* criteria quadrant refer to knowledge base validity, including knowledge representation and associated theories of jurisprudence, inferencing, and the provision of explanations.

Validation of legal modeling by domain legal experts – driven by operational usability assessments – is also mentioned in three methodologies referring to ontological expert knowledge evaluation. For example, the Methodology for Modeling Legal Ontologies (MeLOn) [14] offers evaluation parameters, notably, completeness, correctness, coherence of the conceptualization phase and artifact reusability. Usability was considered in an experimental validation of a legal ontology by legal experts, the Ontology of Professional Judicial Knowledge (OPJK), described in [7]. This model was validated in a two-step process. First, the evaluators answered a questionnaire whereby they expressed their opinion on their level of agreement towards the ontology conceptualization and provided suggestions for the improvement thereof. Then an experimental validation based on a usability questionnaire followed, the System Usability Scale (SUS), tailored to evaluate the understandability and acceptance of the contents of the ontology. This evaluation questionnaire could offer rapid feedback and support towards the establishment of relevant agreement, shareability or quality of content measurements in expert-based ontology evaluation. An evaluation methodology based on Competency Questions (CQs) [18] was built to evaluate the transformation of legal knowledge from a semi-formal form (Semantics Of Business Vocabulary And Rules - Standard English (SBVR-SE)) [13] to a more structured formal representation (OWL 2), and to enable cooperation between legal experts and knowledge IT experts in charge of the modelling in logic formalism.

Although the framework target of this work's analysis (i.e., the DAPRECO Knowledge Base) refers to a validated ontology (i.e., the Privacy Ontology

(PrOnto)), an argument for its legal validity cannot only derive from the validity of the ontology of reference. It requires a more comprehensive analysis and we believe that both qualitative evaluation methodologies and certain criteria from the CCCF are required. Ontologies are in fact about concepts, data, and entities and any validation strategy of them is inevitably about assessing the legal qualities of those objects. Formal models for legal compliance, such as the DAPRECO Knowledge Base, model also the logical and deontic structure of a legal text, its temporal aspects and, as the used formalism yields multiple conflicting interpretations, it includes structural elements to allow defeasible reasoning. The validation assessment should take these elements into account.

Thus, the necessity of an integrated approach, which additionally should also acknowledge an operational-usability assessment, since the legal validity of the DAPRECO Knowledge Base logic formulæ have to be validated by non experts in logic.

3 DAPRECO Knowledge Base

The target of the validation methodology we propose in this work is the so called DAPRECO Knowledge Base. Currently, it contains a preliminary formalization of GDPR's provisions. Technically, the Knowledge Base stands on three interconnected components: legal text; conceptual model; deontic rules. Since it is meant to provide a semi-automated assistance to legal experts, all of the three components need to be machine-readable, and so, consolidated standards and reference formats have been used to model them.

The *legal text* is modelled in Akoma Ntoso². Using ordinary XML parsers, it makes easy to navigate the document and reference specific portions of text. The *conceptual model*, specifically designed using the Web Ontology Language (OWL) language in an XML serialization, is contained in a legal ontology of privacy and data protection concepts, called PrOnto [16,17]³, which the Knowledge Base refers to. The ontology itself, has been developed following the MeLOn methodology, which is based on a glossary and a set of Competency Questions (CQs). The *deontic rules* of the GDPR are expressed in Reified Input/Output (RIO) logic [20]. It is a defeasible deontic logic that uses reification, a technique added to the logic to avoid nested obligations.

This set of RIO formulæ, their consistency and completeness are the real target of the validation task⁴. The formulæ act as a sort of *trait d'union* between the other two components, as they contain references both to ontological elements of the conceptual model and to the textual portions of the legal document expressed in Akoma Ntoso format.

² Currently stored at <https://github.com/guerret/lu.uni.dapreco.parser/blob/master/resources/akn-act-gdpr-full.xml>.

³ Currently stored at <https://github.com/guerret/lu.uni.dapreco.parser/blob/master/resources/pronto-v8.graphml>.

⁴ The formulæ are available at <https://github.com/dapreco/daprecokb>.

All formulæ are if-then rules in the form (x, y) , such that when x is given in input, y is returned in output. When applied to the legal domain, there are three sets to which rules can belong to: C is the set of constitutive norms, which defines when something counts as something else in the domain. Every pair $(x, y) \in C$ reads as “ $x \rightarrow y$ ”, as standard first-order logic implications; O and P are respectively the set of obligations and the set of permissions of the normative system. A pair $(x, y) \in O$ reads as “given x , y is obligatory”, while a pair $(x, y) \in P$ reads as “given x , y is permitted”.

Both the “if” and the “then” part of each formula are composed by a conjunction of predicates. Each predicate is in the form of the predicate name followed by a list of attributes. The name can be a concept belonging to an ontology (e.g., the PrOnto ontology) or it can be a logical operator. For example, $(\text{PrOnto} : \text{PersonalDataProcessing } xz)$ refers to a concept in the PrOnto ontology and takes two arguments. The predicate alone is incomplete, because it also needs to describe the two predicates used as arguments. If x is a controller and z some personal data of a data subject, an example may be formula 1.

$$((\text{prOnto} : \text{Controller } x) \wedge (\text{prOnto} : \text{DataSubject } w) \wedge (\text{prOnto} : \text{PersonalData } z w) \wedge (\text{prOnto} : \text{PersonalDataProcessing}' e_p xz)), \quad (1)$$

Furthermore, in RIO logic a predicate can be *reified* to be used as arguments for other predicates. Thus $(\text{prOnto} : \text{PersonalDataProcessing}' e_p xz)$ is a new predicate, different from $(\text{prOnto} : \text{PersonalDataProcessing}, xz)$; it represents the *possibility* that there is a processing of personal data. This allows e_p to be used as argument to another predicate.

How the DAPRECO Knowledge Base Looks Like. To make RIO formulæ machine-readable format, they were written in LegalRuleML⁵, an XML markup language and a developing OASIS standard for representing the fine-grained semantic contents of legal texts [1]. In essence, each formula expressed as a LegalRuleML rule contains two parts: premise (if) and the consequence (then). The predicates (and their arguments) composing both parts are serialized as RuleML atoms (and variables). The example above, with reification added, is serialized as in Listing 1.

4 Validation Methodology

The object of the validation are the formulæ, regardless of its expressive form (logic or LegalRuleML serialization). The defeasible nature of the logic allows for many interpretations, even one superseding or contrasting with another, as typical in law. Thus, there is no correct interpretation to be validated, rather it is sought the author’s checking whether a logic formula correctly represents one particular interpretation (which can be also his/her own).

⁵ <http://ruleml.org/index.html>.

What we ultimately pursue is a feedback on the *legal quality* of the formulæ’s expressed meaning(s). This quality can be measured at least using metrics, such as: *accuracy* (does the deontic modality expressed by a formula match the corresponding legal provisions? are the relationships among the concept accurately represented?); *completeness* (is all the required domain knowledge explicitly stated, or can it at least be inferred from the vocabulary?); (*subjective*)*correctness* (is the formula’s meaning correct, according to your interpretation?); *consistency* (is the formula’s meaning consistent with the law?); and *conciseness* (is there any amount of redundancy in the representation, or is it concise?).

These metrics can be empirically assessed using an *ad hoc* questionnaire, a very useful quantitative indicator of user acceptance [25]. In this case, where users are lawyers, the questionnaire was designed with the purpose of having legal feedback on the quality of the legal interpretation in the RIO formulæ, and was built around six questions reported below:

q_1	Is the deontic modality (e.g., obligation) of the formula the same as in the article?
q_2	Does the formula capture all the important legal concepts?
q_3	Does the formula capture all the important legal relations?
q_4	Is the interpretation given by the model correct?
q_5	Is the interpretation complete?
q_6	Is the interpretation to the point?

The questions have been tailored to assess Accuracy (q_1); Completeness ($q_2 - q_3$); Correctness (q_4); Consistency (q_5); and Conciseness (q_6).

However, the evaluator needs to understand what a formula states and ideally, the strain to read the formula should not overtake the effort required to provide feedback. From experience gathered in the DAPRECO project we learned that even IT experts required several and repeated explanations to understand what a specific formula expressed. Hence the need for a *human-readable* representation of the formulæ, which preserves the meaning of the machine-readable model but

Listing 1. LegalRuleML of formula 1.

```

<ruleml:Exists>
  <ruleml:Var key=":z">z</ruleml:Var>
  <ruleml:Var key=":x">x</ruleml:Var>
  <ruleml:And>
    <ruleml:Atom>
      <ruleml:Rel iri="prOnto:DataSubject"/>
      <ruleml:Var key=":w">w</ruleml:Var>
    </ruleml:Atom>
    <ruleml:Atom>
      <ruleml:Rel iri="prOnto:PersonalData"/>
      <ruleml:Var keyref=":z"/>
      <ruleml:Var keyref=":w"/>
    </ruleml:Atom>
    <ruleml:Atom>
      <ruleml:Rel iri="prOnto:Controller"/>
      <ruleml:Var key=":x">x</ruleml:Var>
      <ruleml:Var keyref=":z"/>
    </ruleml:Atom>
    <ruleml:Atom keyref=":A3">
      <ruleml:Rel iri="prOnto:PersonalDataProcessing"/>
      <ruleml:Var key=":ep">ep</ruleml:Var>
      <ruleml:Var keyref=":x"/>
      <ruleml:Var keyref=":z"/>
    </ruleml:Atom>
  </ruleml:And>
</ruleml:Exists>

```

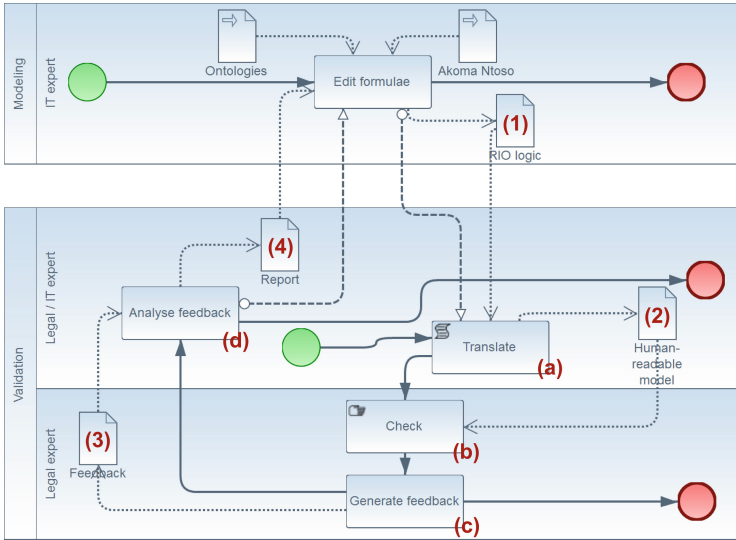


Fig. 1. Workflow of the modeling and of the validation methodology.

is understandable by non-experts in logic, ontologies, or XML. We devised one and we measured its *usability* (what the formula says is easy to understand?).

The methodology workflow is resumed in Fig. 1 on the lower portion of the diagram (“Validation”). The machine-readable version of the modelling of the legal text—in our case, the DAPRECO Knowledge Base—is the output of the modeling effort by the IT expert. That file needs to be processed and rewritten (“Translate”, (a) in Figure) into a human-readable representation. The “Human-readable model” (2) is then validated (“Check”, (b)) against specific measures defining if whether the modeling was correct from a legal point of view. The checking process (“generate feedback”) produces a list of “Feedbacks” (3) expressing the assessment of the model’s legal qualities, likely in the form of quality measures or answers of a questionnaire. The feedback is then analyzed (“Analyse feedback”, (d)), e.g., the statistical significance of certain answers will be measured to compile a “Report” (4) for the IT experts and for the knowledge base builders. The report contains suggestions to review and improve their modelling. This workflow can be iterated until both parties are satisfied.

Due to space constraints, this paper will not delve into the details of each individual step, but only report on the three most critical steps in the methodology: “Translate”, “Check”, and “Generate Feedback”.

4.1 Translate

The “Translate” step generated a representation of the formulæ that legal evaluators could read in order to give feedback about the *legal quality* of the formulæ’s meaning(s). We will refer to this synthetic digest (of an otherwise specific logic formalism) as *human-readable representation of a RIO formula* and herewith we show how it was build and how we measured its understandability.

Translating LegalRuleML of RIO Logic Formulæ. Our input is the DAPRECO knowledge base, a LegalRuleML file of RIO formulæ expressing the legal meaning of articles of the GDPR. Perusal of the knowledge base rendered some difficulties, although slightly facilitated by accompanying comments. For instance, in the LegalRuleML serialization, detecting the enumerated prohibitions, obligations, reparations, exceptions was not straightforward. According to [24] “the list of [LegalRuleML] elements and their definitions are not sufficient for the consistent and accurate application of the annotations to text, nor is there clarification about how to analyse source text into LegalRuleML. Thus, an annotation methodology is required to connect text to LegalRuleML”.

To elicit a set of usability requirements for the human-readable model, we performed an internal unstructured inquiry where legal experts were asked to spell out what was making the reading hardened and mentally burdensome when answering the previous questions. The inquiry highlighted the following obstacles to a clear understanding of the LegalRuleML of a RIO formula: (1) a formula has little structure, and there are many variables and cross-references between them, forcing the reader to move up and down the code; (2) external references may refer to concepts expressed in the PrOnto ontology, or to logical operators from the RIO logic; (3) the choice of the names of predicates and arguments is not driven by a clear strategy, so that the formula appears confusing; (4) whether a formula is an obligation, a permission or an entailment does not immediately stand out from its syntax, as it depends on the context, which is defined elsewhere according to LegalRuleML practices; (5) negations are hard to capture, as they are structured with two predicates, the first introducing the negation of the second predicate that is expressed positively; (6) RIO logic avoids nesting of obligations and permissions, separating the content of the deontic rule from its bearer in two distinct formulæ. This decision, motivated by the purposes of the logic, can create some confusion, as ultimately there will generally be two separate, and almost identical, formulæ, with the same premises and almost the same consequence.

We address all these problems in a two-step “Translation”: the first step is a software that parses the XML, expands and reorders the predicates of the formula; this addresses obstacles 1, 4, 5 and 6. The second is hand-made, to derive an almost natural language break-up version of the formula which, we believe, removes obstacles 2 and 3.

Step One: Automatic Parsing. The output of the automatic translator⁶ overcomes the problems enumerated above in the following way: (i) variables are substituted with the predicate (taken from PrOnto) that restricts their type; (ii) predicates from PrOnto are clearly highlighted in **bold**, whereas predicates from RIO logic and terms that have been introduced for readability’s sake are not; (iii) the translation of a predicate introduces some terms to set everything into context. This technique works quite well due to a good structure of the ontology; (iv) the context of a formula (obligation, permission, constitutive) is carried over to the translation; (v) negations are treated by translating the

⁶ Available at <https://github.com/guerret/lu.uni.dapreco.parser.git>.

predicates in an inline negative sentence. Additionally, when a negation is the object of an obligation, the latter is renamed into a prohibition, and its content expressed positively; **(iv)** if the parser can find another formula with the exact same IF conditions, then they are most likely the content and bearer of an obligation or permission, so the two formulæ are merged into a single translation, which includes both content and bearer.

Article 7.1 of the GDPR can serve as an example: “Where processing is based on consent, the controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data”⁷.

The (simplified) RIO formula that IT experts wrote (and later encoded in LegalRuleML) to model the provision is shown in formula 2.

$$\begin{aligned}
 & \left(\left[(\text{RexistAtTime } a_1 \ t_1) \wedge (\text{and } a_1 \ e_p \ e_{hc} \ e_{au} \ e_{dp}) \wedge (\text{DataSubject } w) \wedge \right. \right. \\
 & \quad (\text{PersonalData } z \ w) \wedge (\text{Controller } y \ z) \wedge (\text{Processor } x) \wedge (\text{nominate}' e_{dp} \ y \ x) \wedge \\
 & \quad (\text{PersonalDataProcessing}' e_p \ x \ z) \wedge (\text{Purpose } e_{pu}) \wedge (\text{isBasedOn } e_p \ e_{pu}) \wedge \\
 & \quad \left. (\text{Consent } c) \wedge (\text{GiveConsent}' e_{hc} \ w \ c) \wedge (\text{AuthorizedBy}' e_{au} \ e_{pu} \ c) \right] \rightarrow \\
 & \quad \left[(\text{RexistAtTime } e_a \ t_1) \wedge (\text{AbleTo}' e_a \ y \ e_d) \wedge (\text{Demonstrate}' e_d \ y \ e_{hc}) \right] \in \mathcal{O} \quad (2)
 \end{aligned}$$

The parser translates the formula as follows:

<p>IF, in at least a situation,</p> <ul style="list-style-type: none"> - At time $:t_1$, the following situation exists: <ul style="list-style-type: none"> • (All of the following $(:a_1)$) <ol style="list-style-type: none"> 1. Processor $(:x)$ does PersonalDataProcessing $(:e_p)$ of PersonalData $(:z)$ 2. DataSubject $(:w)$ performs a GiveConsent $(:e_{hc})$ action on Consent $(:c)$ 3. Purpose $(:e_{pu})$ is AuthorizedBy $(:e_{au})$ Consent $(:c)$ 4. Controller $(:y)$ nominate $(:e_{dp})$ Processor $(:x)$ - PersonalData $(:z)$ is relating to DataSubject $(:w)$ - The Controller $(:y)$ is controlling PersonalData $(:z)$ - PersonalDataProcessing $(:e_p)$ isBasedOn Purpose $(:e_{pu})$ <p>THEN it must happen that, in at least a situation,</p> <ul style="list-style-type: none"> - At time $:t_1$, Controller $(:y)$ is Obligated to AbleTo $(:e_a)$ - Controller $(:y)$ Demonstrate $(:e_d)$ GiveConsent $(:e_{hc})$
--

Although the translation still requires some mental effort to be processed, it is at least understandable without having expertise in logic. The automatic processing also allowed the modeller to verify that the intended meaning has not been changed and is preserved in the translation.

Step Two: Hand Made Break-Up. The automatic translation has been further hand-processed. The output is a natural language break-up that highlights the following elements: Premises and the Conclusion of the formula; the Deontic Modality, the Ontological Concepts that can be recognized in the article, Other Ontological Concepts present in the formula but not mentioned in the article; the Contextual meaning, which is what the formula expresses but is not in the article, and the Overall Meaning of the formula. The break-up of Article 7.1 is shown in Table 1.

⁷ The full translations for Articles 5.1 and 7.1 can be found in the repository from note 6, in the “jurisin” folder.

Table 1. Structure of the formula’s meaning.


Premise	Where processing is based on consent,
Conclusion	The controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data
Modality	Obligation
Ont. Concepts	Where [Processing] is based on [Consent], the [Controller] shall be [Able to] [Demonstrate] that the [Data subject] [Has consented = GiveConsent] to [Processing] of his or her [Personal data]
Other Ont. Concepts	[Purpose]; [Processor]; [IsAuthorizedBy]; [Nominates]; [IsBasedOn]; [BeAbleTo]
Context	There is a processing, which has a purpose authorized by a consent given by a data subject, and that is what a processor, whom a controller controlling the personal data nominates, does on personal data of the data of the data subject
Overall Meaning	Whenever there is a processing, which has a purpose authorized by a consent given by a data subject, and that is what a processor, whom a controller controlling the personal data nominates, does on personal data of the data of the data subject then the controller is obliged to able to demonstrate that “data subject gave consent”

Measuring the Usability of the Human-Readable Model. Before collecting feedback on the quality of the model, the human-readable model must be able to be read consistently and correctly by evaluators. Hence, our experiment consisted in requesting four legal evaluators (two with knowledge of deontic logic, two without it) to answer a few yes/no questions about their understanding of the models of two GDPR provisions. Our priority was to check the modeling of different types of legal norms into logical formulae and as such, Article 5.1(a) was elicited as it represents a constitutive rule, and Article 7.1 evinces an obligation. The input is the human-readable model, but we also fed the original XML formalization and the pre-processed output as control cases, measuring (pure, not Fleiss Kappa) the average interrater agreement between the answers of the evaluators for each model. The questions, built in the wake of the ones used for the validation check (the initial questionnaire) were the following: 1. *Can you identify the formula’s premise?* 2. *Can you identify the formula’s conclusion(s)?* 3. *Can you identify the deontic modality (obligation, permission, other)?* 4. *Can you identify the formula’s explicit ontological concepts?* 5. *Can you identify the formula’s implicit ontological concepts?* 6. *Do you understand what the formula means?* 7. *Try to rewrite the formula in your own words. Did you succeed?*

We measured the average agreement over all questions and the two formulae. The agreement on the answer ‘yes’, indicating readability, are shown in Table 2. The hand-processed model is where the evaluators, including the laymen in logic, agree almost unanimously over answering ‘yes’ to all questions, thus indicating high understandability; the control item, the XML file, is where instead there is a major consensus on not being understandable. Our result also reflects that validators already knowledgeable on logic can somehow read the XML files, despite

Table 2. Output of the agreement ‘yes’ on the readability experiment.

	Commented XML	Intermediate	Human-readable
all	39.4% (yes)	40.9%(yes)	97.7%(yes)
lay	0% (yes)	45.5%(yes)	95.5%(yes)



not fully; unsurprisingly, non-experts thereof could not make any sense of it. Conversely, there is no consensus on the understandability of the automatically-processed model. Supposedly, better usability scores may be attained by training the legal evaluators, but we have not explored this possibility. In this particular experiment, we did evaluate other qualities of the model, such as its correctness. Correctness has been assessed in a second experiment, see next section.

4.2 Check and Generate Feedback

We measured understandability as the inter-tester agreement: this measure can suffice to the present goal of having the human-readable model as a candidate within the methodology, although additional measures can provide a deeper evaluation of its usability. More evidence would be needed to assert that our hand-processed model is readable, but since our evaluators generally agreed on its understandability, it can already be used to collect answers to questions $q_1 - q_6$. This is what we did as next steps in the methodology, together with the analysis of the feedback collected during this research.

The starting point is the human-readable representation of Articles 5.1(a) and 7.1 of the GDPR. The ‘‘Check’’ action (see Fig. 1) has been implemented by gathering a set of four validators, all jurists knowledgeable on data protection law, and by asking them to answer questions $q_1 - q_6$ of the questionnaire.

Evaluators were told to compare the meaning of the formulæ, as expressed in the human-readable representation of the RIO logic, with the legal interpretation

that they would convey to the articles of the GDPR. We also (re)-asked them a few questions meant to reveal how much understandable for them is the human-readable format, before they start using it. General understandability of the format was assessed already, but here the assessment is meant as a trust measure over the expert’s answers. From those trusted answers, we therefore compiled a few recommendations. This is the “Generate Feedback” step in Fig. 1.

While the evaluators were requested to answer the questionnaire in reference to each of the three expressions of the formula (logic, automated translation, and manual break-up), the results are shown for brevity’s sake only for the final format. Feedbacks on the less-readable formats have been used to refine the two steps of the translation. Additionally, the multiple feedbacks helped detect the exact location of errors, whether in the formula, in the automated translation, or in the manual break-up.


Questions $q_1 - q_6$ are yes/no questions but we invited our checkers to motivate the answers and to pinpoint whatever observation they valued meaningful. We collected eight documents (four reviewers, two articles) with such written answers and comments which we reviewed and summarized. The following table resumes the findings, wherein we report the comments whenever the answer to the question was ‘no’, indicating that someone found some issue pertinent.

Table 3. Feedback collected

	Art 5.1a	Art. 7.1
Accuracy	✓	✓
Completeness	It was complex to capture the legal concepts within the structure of the formula; It is missing the obligation: “the processing must be fair, lawful, transparent”	It was complex to capture the legal concepts within the structure of the formula;
Consistency	Interchanged roles for the controller and the processor; The interpretation is complex. It refers to the implementation and description of a measure that it is hard to understand; I can read/understand the model, but I think it does not faithful to the article’s meaning;	The reference to consent should be enhanced, namely regarding the requirements concerning the burden of the proof; “Shall” is not captured;
Conciseness	The formula mentions “implement” “describe” not expressed in the article; “implement measure” is not expressed in the article; “Obligated to be able” sounds weird;	It is redundant and restates concepts already present at previous articles;

Table 3 shows that legal experts were able to give feedback on all the factors about the quality of the legal interpretation in the logic formalization of the articles. Even if the input to provide to the IT expert is not yet straightforward, a few highlights clearly emerge.

For instance, all experts easily understood and confirmed the deontic modality and agreed that the formulæ captured all the legal concepts and relations (see Table 4). But is from the analysis conferred to the provided comments that we are able to offer a broader spectrum, for they refer to the above surveyed

Table 4. Inter-evaluators agreement on answering ‘yes’ to the questions

criteria and also to other (non-surveyed) related criteria. Comments – in Completeness like “it was complex to capture the legal concepts within the structure of the formula”; comments in Consistency like “It refers to the implementation and description of a measure that it is hard to understand; “It is redundant and restates concepts already present at previous articles”, and comments in Conciseness like “‘Obligated to be able to’ sounds weirds” – clearly show uneasiness about how formula have been structured; such comments may lead to a better formalization, for instance, stating certain contextual facts as a common premise valid for all the GDPR’s articles without repeating them each time.

One evaluator, in particular, has mentioned “Interchanged roles for the controller and the processor” in Consistency. Even if that is stated in the context of the human-readable table, the evaluator was probably induced in error/confused by the excess of information provided. Further analysis is of course required. Extracting from the non-structured comments valid input for the IT expert has to be left as future work, as we comment in the following section.

5 Conclusions and Future Work

This paper leverages a methodology that advocates an interdisciplinary validation of a representation of the GDPR articles in a logic formalism (i.e., RIO logic) to pursue quality, accountability, and transparency within. One important output of the methodology is the production of feedback derived from the involvement of legal experts, while assessing the quality of the legal interpretation that IT experts may instill in the formalization of the GDPR. This work has gathered evidence that such step is feasible. As a proof-of-concept, a small number of legal experts has been asked to answer six questions with the purpose of collecting comments about how two logic formulæ, modelling Articles 5.1a and

7.1 of the GDPR, are complete, accurate, concise, and consistent in reflecting the legal meaning of the articles. Several comments have been collected. Although a thorough analysis thereof requires more time – an involvement of a larger group of expert checkers is also advisable – we were able to identify a few issues of relevance which the IT expert can account in the formalization work.

Several challenges await us in the near future. We need to improve scalability in producing a human-readable representation of the RIO formulæ: it is currently done manually, starting from the pre-processed version. This is already more readable than the original LegalRuleML version and warrants us that the work to produce a natural language analysis break-up table can be automatized. This step done, a forth bringing process will consist in streamlining the validation of the RIO formalization of the GDPR as a whole. This likely requires to set up an application where the modeling of the IT expert can be suitably translated into the human-readable format and displayed, for online checking, to a group of legal testers in order to provide feedback, until a good assessment of the legal interpretations is reached.

Concomitantly, there is a need to define, together with the legal experts, a more complete set of qualities and possibly a few metrics, which we can quantify and define criteria on the legal quality of the formalization. In Sect. 2 we pointed out possible metrics, and in this paper we have assessed a few (completeness, consistency, conciseness in Sect. 4), but a wide and systematic investigation of the state-of-the-art in this topic has not been done yet. The quadrant criteria presented in [10] also merits attention. This may lead to a revision of the current human-readable model.

References

1. 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/10.1007/978-3-319-21768-0_6
2. Bartolini, C., Giurghi, A., Lenzini, G., Robaldo, L.: Towards legal compliance by correlating standards and laws with a semi-automated methodology. In: Bosse, T., Bredeweg, B. (eds.) BNAIC 2016. CCIS, vol. 765, pp. 47–62. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67468-1_4
3. Bartolini, C., Lenzini, G., Santos, C.: A legal validation of a formal representation of GDPR articles. In: Proceedings of the 2nd JURIX Workshop on Technologies for Regulatory Compliance (Terecom) (2018)
4. Bartolini, C., Lenzini, G., Santos, C.: An interdisciplinary methodology to validate formal representations of legal text applied to the GDPR. In: Proceedings of the Twelfth International Workshop on Juris-Informatics (JURISIN), November 2018
5. Berman, D.H.: Developer’s choice in the legal domain. In: Proceedings of the Third International Conference on Artificial Intelligence and Law (ICAIL). ACM, June 1991
6. Boella, G., Humphreys, L., Muthuri, R., Rossi, P., van der Torre, L.W.N.: A critical analysis of legal requirements engineering from the perspective of legal practice. In: IEEE 7th International Workshop on Requirements Engineering and Law (RELAW), pp. 14–21. IEEE (2014)

7. Casellas, N.: Ontology evaluation through usability measures. In: Meersman, R., Herrero, P., Dillon, T. (eds.) OTM 2009. LNCS, vol. 5872, pp. 594–603. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-05290-3_73
8. Conrad, J.G., Zeleznikow, J.: The significance of evaluation in AI and law. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law (ICAIL), pp. 186–191. ACM, June 2013
9. Conrad, J.G., Zeleznikow, J.: The role of evaluation in AI and law. In: Proceedings of the Fifteenth International Conference on Artificial Intelligence and Law (ICAIL), pp. 181–186. ACM, June 2015
10. Hall, M.J.J., Hall, R., Zeleznikow, J.: A process for evaluating legal knowledge-based systems based upon the Context Criteria Contingency-guidelines Framework. In: Proceedings of the Ninth International Conference on Artificial Intelligence and Law (ICAIL), pp. 274–283. ACM, June 2003
11. Hall, M.J.J., Zeleznikow, J.: Acknowledging insufficiency in the evaluation of legal knowledge-based systems. In: Proceedings of the Eighth International Conference on Artificial Intelligence and Law (ICAIL), pp. 147–156. ACM, May 2001
12. Koers, A.W.: Knowledge Based Systems in Law, 1st edn. Kluwer Law and Taxation Publishers, Deventer (1989)
13. Lévy, F., Nazarenko, A.: Formalization of natural language regulations through SBVR structured English. In: Morgenstern, L., Stefaneas, P., Lévy, F., Wyner, A., Paschke, A. (eds.) RuleML 2013. LNCS, vol. 8035, pp. 19–33. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39617-5_5
14. Mockus, M., Palmirani, M.: Legal ontology for open government data mashups. In: Parycek, P., Edelmann, N. (eds.) Proceedings of the 7th International Conference for E-Democracy and Open Government (CeDEM), pp. 113–124. IEEE Computer Society, May 2017
15. Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., Robaldo, L.: Legal ontology for modelling GDPR concepts and norms. In: Proceedings of the 31st International Conference on Legal Knowledge and Information Systems (JURIX), December 2018 (forthcoming)
16. Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., Robaldo, L.: PrOnto: privacy ontology for legal compliance. In: Proceedings of the 18th European Conference on Digital Government (ECDG), October 2018 (upcoming)
17. Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., Robaldo, L.: PrOnto: privacy ontology for legal reasoning. In: Kö, A., Francesconi, E. (eds.) EGOVIS 2018. LNCS, vol. 11032, pp. 139–152. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98349-3_11
18. Ramakrishna, S., Górski, L., Paschke, A.: A dialogue between a lawyer and computer scientist. *Appl. Artif. Intell.* **30**(3), 216–232 (2016)
19. Reich, Y.: Measuring the value of knowledge. *Int. J. Hum. Comput. Stud.* **42**(1), 3–30 (1995)
20. Robaldo, L., Sun, X.: Reified Input/Output logic: combining Input/Output logic and reification to represent norms coming from existing legislation. *J. Log. Comput.* **27**(8), 2471–2503 (2017)
21. Sergot, M.: The Representation of Law in Computer Programs. The A.P.I.C. Series, Chap. 1, vol. 36, pp. 3–67. Academic Press, Cambridge (1991)
22. Stranieri, A., Zeleznikow, J.: The evaluation of legal knowledge based systems. In: Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL), pp. 18–24. ACM, June 1999
23. Susskind, R.E.: Expert systems in law. Out of the research laboratory and in the marketplace. In: Proceedings of ICAIL-1987, Boston, MA, pp. 1–8. ACM (1987)

24. Wyner, A., Gough, F., Levy, F., Lynch, M., Nazarenko, A.: On annotation of the textual contents of Scottish legal instruments. In: Proceedings of the 30th International Conference on Legal Knowledge and Information Systems (JURIX), vol. 302, pp. 101–106. IOS Press, December 2017
25. Zeleznikow, J.: The split-up project. *Law Probab. Risk* **3**(2), 147–168 (2004)