

Towards legal compliance by correlating Standards and Laws with a semi-automated methodology¹

Cesare Bartolini^a

Gabriele Lenzini^a

Livio Robaldo^a

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

Abstract

Since legal regulations do not generally provide clear parameters to determine when their requirements are met, achieving legal compliance is not trivial. If there were a clear correspondence between the provisions of a specific standard and the regulation's requirements, one could implement the standard to claim a presumption of compliance. However, finding those correspondences is a complex process; additionally, correlations may be overridden in time, for instance, because newer court decisions change the interpretation of certain provisions. To help solve this problem, we present a framework that supports legal experts in recognizing correlations between provisions in a standard and requirements in a given law. The framework relies on state-of-the-art Natural Language Semantics techniques to process the linguistic terms of the two documents, and maintains a knowledge base of the logic representations of the terms, together with their defeasible correlations, both formal and substantive. An application of the framework is shown by comparing a provision of the European General Data Protection Regulation against the ISO/IEC 27018:2014 standard.

1 Introduction

Generally, achieving legal compliance is not a trivial task for enterprises. Laws and regulations generally do not specify what measures should be implemented to match the requirements that they state. For instance, in the European Union (EU), where there is the need to support a single market while letting each country establish its own legal framework, legislation normally defines the safety/security of products and systems only at a very abstract level, leaving to the Member States the choice on how to demonstrate compliance².

In this situation, implementing standards is a viable way for enterprises to create a “presumption of conformity with the specific legal provision they address” [6]. Implementing a standard *per se* does not guarantee legal compliance, but it can provide a significant clue of conformity with regulations. When widely adopted and subject to repeated audits by conformity-assessing bodies, a standard offers an argument of compliance. Of course, such an argument is a presumption, giving a plaintiff the possibility to demonstrate that the organization failed to comply with the legal framework. Still, this *inversion of the burden of proof* is often preferable to a personalized solution [6]. However, the problem of establishing and assessing this presumption remains, and it may become a serious issue when new laws reshape the legal landscape for businesses: while awaiting for the establishment of feasible solutions, an enterprise faces the risk of liability for not being compliant and may, on that ground, be sanctioned.

¹Livio Robaldo has received funding from the European Unions H2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 661007. Cesare Bartolini has received funding from the European Union's H2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 690974.

²Regulation (EU) 1025/2012, Article 2.1

A domain that will need to be thoroughly addressed from this point of view is data protection, because the recent adoption of the General Data Protection Regulation (GDPR)³, which will be applied from 25 May 2018⁴, will pose significant challenges for undertakings in terms of ensuring compliance with it, as it brings major changes to the regulatory framework for personal data protection in Europe. On the other hand, the ISO/IEC 27018:2014 standard, concerning public clouds acting as personal data processors, can be regarded as a building block [6] that helps data-processing organizations comply with the principle of accountability and with their obligations resulting from data protection laws.

We propose a software framework that aids in determining the formal and substantive correlations between the provisions in a standard and in a law. The framework’s core is a logic-based methodology to represent, in a machine-processable format, (a) the relevant syntactic concepts in the provisions, and (b) the relevant correlations between them. In this paper, we describe this logic-based methodology and exemplify how it works using provisions from the ISO 27018 standard and the GDPR.

The framework depends on two auxiliary functional blocks (see Section 3): (i) a *logic knowledge base* that can be used, corrected and extended by legal experts, that stores the machine-processable logic correlations; (ii) a *set of Natural Language Semantics (NLS) and Natural Language Processing (NLP) techniques* that allow a user to browse a XML representation of the documents and to search and retrieve the words, terms, and sentences that have been found relevant for correlation. The NLS and NLP techniques will help users show the established correlations within the knowledge base. Any expert user can contribute to reinforce, correct, justify, and expand the correlations. Due to differences in legal interpretation, some of the correlations could be in contradiction. Interpretations from authoritative sources (such as high courts) will eventually sort contradictions out.

Technically, the different correlations are expressed in a deontic and defeasible logic for legal semantics called Reified Input/Output Logic (see Section 2). The selection of relevant terms and the definition of correlations, requiring human reading, processing and decision-making, is therefore semi-automatic. It will be supported by the extensible knowledge base; the process of browsing, aided by the tools and techniques from the NLS and NLP domains, add efficiency and precision.

2 Related work

Reified Input/Output Logic is the defeasible logic that we propose as the formal language to express correlations. It was designed as an attempt to extend the state-of-the-art research in legal informatics, by investigating the *logical architecture* of the provisions, which are available in natural language only.

Reified Input/Output logic is a logical framework [15] for representing the meaning of provisions. Contrary to the large majority of its competitors (e.g., [8, 10, 7]), Reified Input/Output logic integrates modern insights from the NLS literature. Specifically, it merges Input/Output logic [11], a well-known formalism in Deontic Logic, with a first-order logic for NLS grounded on the concept of *reification* [9]. Reification [4] is a concept that allows to move from standard notations in first-order logic such as “(give $a b c$)”, asserting that “ a ” gives “ b ” to “ c ”, to another notation in first-order logic “(give $e a b c$)”, where “ e ” is the *reification* of the giving action. “ e ” is a first-order logic term denoting the giving event of “ b ” by “ a ” to “ c ”. Thanks to reification, the logic is able to express a wide range of phenomena in NLS such as named entities, quantification, anaphora, causality, modality, time, and others. In particular, the simplified version of Reified Input/Output logic that we use in our example is an extension of first-order logic that distinguishes three kinds of implication: “ \rightarrow ”, “ \rightsquigarrow ”, and “ \Rightarrow ”.

The implication “ \rightarrow ” is the standard trust-value implication of first-order logic, whereas “ \rightsquigarrow ” is its *defeasible* [14] version. Defeasible here means that an implication “ $\Phi \rightsquigarrow \Psi$ ” holds by default unless overridden by “stronger” implications. When instantiated properly, this notion of stronger implications resolves the potential contradictions emerging because of the non-monotonic nature of the defeasible reasoning. Reified Input/Output logic also includes other mechanisms to deal with unresolvable conflicts⁵. A possible solution to deal with this type of conflicts is to leave the conflict open until more evidence will allow the reasoner to take a decision. This is, in short, what better fits a situation with conflicts due to multiple legal interpretations. The third implication, “ \Rightarrow ”, is a deontic implication. Taken a

³Regulation (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).

⁴GDPR, Article 99.2.

⁵See plato.stanford.edu/entries/logic-nonmonotonic/#sec-2-2.

formula Φ , referring to a set of pre-conditions, and another formula Ψ , referring to a set of actions, the meaning of “ $\Phi \Rightarrow \Psi$ ” is “given the pre-condition Φ , actions Ψ are obligatory”, that is, the actions must be undertaken if the pre-conditions hold. Note that, according to the interpretation of “ \Rightarrow ”, the formula “ $\Phi \Rightarrow \Psi \wedge \Phi \wedge \neg\Psi$ ” is not inconsistent: it only means obligation Ψ has been violated.

The framework in [15] further distinguishes the formulæ belonging to the assertive contextual statements (ABox), *i.e.*, the formulæ denoted by the provisions, from those belonging to the terminological declarative statements (TBox), *i.e.*, the definitions, axioms, and constraints on the predicates used in the ABox formulæ. Formulæ in the ABox are in the form “ $\forall x_1 \dots \forall x_n [\Phi(x_1 \dots x_n) \Rightarrow \Psi(x_1 \dots x_n)]$ ”, where arguments “ $\Phi(x_1 \dots x_n)$ ” and “ $\Psi(x_1 \dots x_n)$ ” are conjunctions in standard first-order logic. Formulæ in the TBox can be any formula in standard first-order logic augmented with the operator “ \rightsquigarrow ”.

The GDPR is the most recent step in the evolution of data protection rules in the European Union. The Regulation, which will apply from May 2018, replaces the current Directive 95/46/EC⁶, which introduced a minimum level of protection but required implementation by means of Member State laws. As such, the Directive produced a heterogeneous legislation throughout the European Union, although with a common base of protection. On the other hand, the GDPR, being a Regulation, will be directly applicable in all Member States, and thus create a homogeneous data protection framework in the EU.

The purposes of the GDPR, according to its promoter [13], are to align data protection rules with the most recent developments in data-processing technologies, while still providing a legislation that is flexible enough not to become outdated over the course of a few years. In addition, the rights of the data subject are strengthened by burdening the data controller with new obligations, and enforcing such obligations with heavy penalties. Controllers and processors are required to be compliant with the Regulation, which sets up fines as high as four percent of the total annual worldwide turnover of a company in some cases of infringements of its provisions.

The interpretation of the Regulation will be provided mainly by doctrine and jurisprudence. The latter is not available yet: since the Regulation was recently published and won’t be applicable for the next two years, no decisions based on its provisions exist yet. When it will be applied, relevant decisions on its interpretations are expected to be issued by the Data Protection Authorities (DPAs) of Member States, by national courts, and by the Court of Justice of the European Union (CJEU). On the other hand, some doctrinal analysis from legal experts already exists (*e.g.*, [5]), and significant literature will be published in the upcoming months.

3 Methodology

The framework we propose offers a computer-aided methodology to analyze standards to make an argument of compliance with respect to a specific piece of legal text, and is schematically summarized in Figure 1. Users (who may be lawyers, regulators, auditors, or other legal experts) access a digital and annotated XML representation of the normative texts (laws and standards). While browsing a document and selecting the relevant concepts, NLP and NLS tools help traverse the rest of the documents, find related terms, and recall previous correlations between them. Correlations from different sources have different degrees of importance, which need to be tracked using specific metadata. The framework implements a collaborative strategy to evaluate the stored correlations. The user’s decisions are stored in the knowledge base, after being appropriately represented in a logic for legal semantics.

The framework, and in particular its knowledge base, does not pretend to be complete. Rather, it provides the expert user with an updated knowledge that helps him take autonomous and informed decisions, both when confirming the correlations the tool suggests and when choosing to define new correlations. The knowledge base is designed to support defeasible reasoning, *i.e.*, to tolerate (apparent) inconsistencies of different interpretations of terms, by overriding general assertions into more contextually-specific ones. Conflicts are especially frequent in legal interpretation, but they can generally be solved considering that the interpretation by higher-instance courts, such as the CJEU, prevails over lower ones. In order to cope with interpretations of different legal weights, which may supersede one another, the logic formalism that the framework embeds is defeasible: correlations can be updated, modified, rewritten and weighed. If conflicts do remain, the framework still embeds strategies that help the user take a decision.

⁶Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.

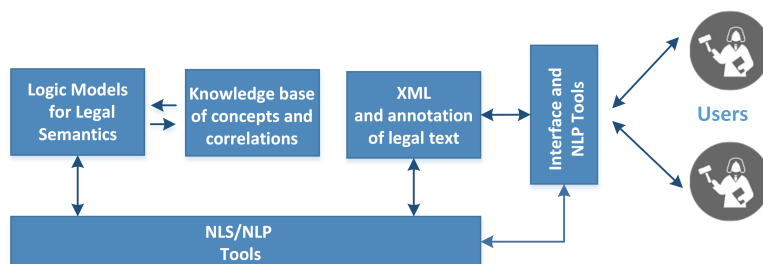


Figure 1: The framework at a glance.

The methodology we propose is highly interdisciplinary, involving a strict interaction between law and computer science. Its purpose is to establish correlations between the provisions in the standards and those in the law. As pointed out in Section 1, correlations can be further divided into two different categories: *formal* and *substantive* correlations.

Formal correlations entail a mere textual overlap between concepts. For example, formal correlations would allow us to observe that both the GDPR and the ISO 27018 standard use the term “notify” (see Section 4). On the other hand, substantive correlations are more complex and entail the analysis of the actual meaning of terms. To assert a correlation of this kind, requirements must be met in a concrete way. Following the previous example, to assert a correlation between a provision of the GDPR and one of the standard concerning notification, it is necessary to verify the exact meaning of the term “notify” in the two texts.

Implications that define a match are *defeasible* and can be overridden, for instance, because of a new decision by a court or DPA, or simply because of the evolution of the interpretation in doctrinal analysis. This way, the output of the methodology can be easily kept up to date by introducing new correlations as they are developed by legal actors.

The methodology follows three steps to build the correlations, involving both a legal and a technical approach. The legal approach is focused on the interpretation of the provisions of laws (the GDPR in our example) and standards (ISO 27018), whereas the technical approach consists of modelling those provisions into an ontology, and expressing the interpretation by means of logical formulæ.

Step 1: analysis of the provisions. The provisions of the law and of the security standard are analyzed by legal experts, who provide an interpretation of the terms used in the provisions and compare them in search of semantic correlations. There is no need for this interpretation to be final, as more interpretations can be added later, and old interpretations can be overridden by newer ones, but this start requires a significant manual activity.

To support the execution of this step, legal experts are assisted by external NLP procedures that suggest (semi-automatically and during the browsing of the documents) previous translations and correlations on the basis of the information currently stored in the knowledge base. Ultimately, it is the legal expert who must decide whether and how correlations need to be overridden. In that case, new correlations are added and annotated with the source which contributed to define it (*e.g.*, Court of Justice of the European Union, *Dapreco and Copreda Corp.*, C-XYZ/16). Implications by more authoritative sources override defeasible implications by less authoritative ones.

Step 2: creation of legal ontologies. The legal interpretations are mapped onto legal ontologies of the law and the security standard. Legal ontologies [2] model the legal concepts, parties and stakeholders affected by the law, the duties and rights of each stakeholder, and the sanctions for violating the duties. As per ontologies in general, legal ontologies are expressed in a knowledge representation language. For this work, we chose the popular abstract language OWL, which can be serialized using various XML notations. For example, the OWL representation of the data protection ontology will contain concepts such as “controller”, “data subject”, “personal data”, “processing” and so on.

A preliminary version of a legal ontology for the GDPR has been defined already [1]. Albeit partial and based on an older version of the GDPR, it was designed to express the duties of the controller. As such, it can be used to find the correspondences between the requirements expressed in the GDPR and in security standards, until an improved ontology is built.

Step 3: generation of logic formulæ. The third and final step of the methodology consists of generating the logical formulæ representing the set of provisions in the law and the set of provisions in the security standard, as well as the implications between them. These formulæ are expressed in Reified Input/Output logic [15]. An example is shown in the next section.

Associating textual provisions to logical formulæ amounts to converting ambiguous and vague terms into non-ambiguous items (predicates and terms). Words in the provisions are represented via predicates reflecting their vagueness. For example, the word “notify”, included in the sample provisions used in Section 4, will be represented via the homonym predicate “notify”. These “vague” predicates may be defined by adding implications and further constraints (axioms). Those implications will be *defeasible*, so that they can account for different legal interpretations. Predicates are associated with classes of the ontologies developed in Step2 or with standard general-purpose ontologies/repositories belonging to the NLS literature, e.g. Verbnnet [16].

4 Generation of Logic Formulæ: example

We exemplify step 3 of our methodology. This step, which lies at the core of the methodology, is the most technical, and more innovative than steps 1 and 2 which instead rely upon existing techniques. We use a provision from the GDPR and an article of the ISO 27018 security standard:

- (a) GDPR, Article 33.2: *The processor shall notify the controller without undue delay after becoming aware of a personal data breach.*
- (b) ISO 27018, Article 9.1: *The public cloud PII processor should promptly notify the relevant cloud service customer in the event of any unauthorized access to PII.*

The formalization of the two provisions in the simplified version of Reified Input/Output logic we use in this paper is rather intuitive. As explained above in Section 3, the formulæ will include predicates reflecting the vagueness of the terms occurring in the sentences. Thus, for instance, the verb “notify”, which occurs in both provisions, is formalized into an homonym predicate “notify”.

On the other hand, the provisions in the ISO 27018 use the term “PII” (Personally Identifiable Information) while the ones in the GDPR use the term “personal data”. Although it sounds rather obvious to consider the two terms as synonyms (as suggested in ISO 27018, Article 0.1), and thus associate them with the same predicate, our methodology keeps them distinct, *i.e.*, it formalizes them via two different predicates “personalData” and “PII”. An additional axiom is then added to the TBox, in order to correlate the two predicates:

$$\forall x[(PII\ x) \rightsquigarrow (personalData\ x)] \quad (1)$$

The axiom correlates the two predicates via a defeasible implication: all that is considered PII is also by default considered personal data. Note that (1) allows the knowledge base to include instances of PII which are *not* personal data. Those would be *exceptions* to the default rule in (1), and it would be necessary to add further higher-priority axioms in order to consistently account for them.

In the light of this, the GDPR provision in (a) is formalized as follows:

$$\begin{aligned} & \forall_e \forall_x \forall_y \forall_z \forall_{e_p} \forall_{e_c} \forall_{e_b} [((dataProcessor\ x) \wedge (dataController\ y) \wedge (personalData\ z) \wedge \\ & \quad (process\ e_p\ x\ z) \wedge (control\ e_c\ y\ z) \wedge (dataBreach\ e_b\ z)) \\ & \quad \Rightarrow \exists_{e_n} [(notify\ e_n\ x\ y\ e_b) \wedge (nonDelayed\ e_n)]] \end{aligned} \quad (2)$$

In (2), “ e_p ”, “ e_c ”, and “ e_b ” are variables referring to three events. “ e_p ” is an event of processing the personal data “ z ” (patient) performed by the data processor “ x ” (agent). “ e_c ” is an event of controlling the personal data “ z ” (patient) performed by the data controller “ y ” (agent). Finally, “ e_b ” is an event of data breach of the personal data “ z ” (patient). Note that the agent of the data breach has been omitted; that could be unknown, but this is irrelevant for the obligation in (2), which concerns the data processor.

(2) states that for any data breach of personal information “ y ”, it is mandatory for the data processor “ x ” who processes “ y ” to notify the data controller “ z ” who controls “ y ”. Such a notifying event “ e_n ” must be carried out without undue delay; this is formalized via a unary predicate “nonDelayed”, which is assumed to be true if the event in its argument has been performed without undue delay.

In (2), “notify” and “nonDelayed” are predicates whose meaning is subject to different legal interpretations. Recalling the difference between *formal* and *substantive* compliance outlined in Section 1, we note that the formalization in (2) only enforces formal compliance. The formula in (2) simply requires the data controller to notify data breaches without undue delay, but it does not specify *how* notifications should be performed for being legitimate.

For instance, the data processor could require the data controller to acknowledge the notification, in order to make sure it was received. Similarly, the processor could be required to avoid sending notifications of data breaches via standard paper mail, in that the time needed by the postal service to deliver the mail could be considered as an undue delay. It is up to judicial authorities to establish the substantive compliance of the obligation in (2).

Of course, we do not have the authority to decide whether (2) is performed in the proper way. In our work, we only aim at providing a methodology to keep track of all legal interpretations of the provisions. From a formal point of view, the knowledge base must be enriched with axioms defining the conditions under which the predicates in the formula are true. Note that different authorities could establish different (conflicting) interpretations of these predicates; this is why *defeasible* implications are needed to model their substantive meaning.

Specifically, the TBox of the ontology will include defeasible axioms that define when, by default, a data processor properly notifies a data controller. For instance, by assuming that email with electronic signature and registered high-priority paper mail are both proper and prompt means to notify the data controller, the following two (defeasible) axioms are added to the TBOX.

$$\begin{aligned} \forall_x \forall_y \forall_{e_1} \forall_{e_2} [(\text{sendEmailWithES } e_1 \ x \ y \ e_2) \rightsquigarrow ((\text{notify } e_n \ x \ y \ e_2) \wedge (\text{nonDelayed } e_n))], \\ \forall_x \forall_y \forall_{e_1} \forall_{e_2} [(\text{sendRegHPMail } e_1 \ x \ y \ e_2) \rightsquigarrow ((\text{notify } e_n \ x \ y \ e_2) \wedge (\text{nonDelayed } e_n))] \end{aligned} \quad (3)$$

In case a judicial authority later decides, for example, that emails with electronic signature are no longer proper and prompt means for notifying data breaches, an additional higher-priority axiom will be added to the TBox in order to override the first axiom in (3). Axioms may be associated to time stamps (although this is not shown in (3)), so that the new axiom will override the old one only for notifying actions performed after a certain date, *i.e.*, since the new decision was issued. On the other hand, the old axiom will still assert that emails with electronic signature are proper and prompt means for all notifications performed before that date.

Formula (4) models the ISO 27018 provision in (b):

$$\begin{aligned} \forall_e \forall_x \forall_y \forall_z \forall_{e_p} \forall_{e_c} \forall_{e_a} [((\text{PIIProcessor } x) \wedge (\text{PIIController } y) \wedge (\text{PII } z) \wedge \\ (\text{process } e_p \ x \ z) \wedge (\text{control } e_c \ y \ z) \wedge (\text{unauthorizedAccess } e_a \ z)) \\ \Rightarrow \exists_{e_n} [(\text{notify } e_n \ x \ y \ e_a) \wedge (\text{promptly } e_n)]] \end{aligned} \quad (4)$$

As it was done for formalizing (a) into (2), the formula introduces predicates that reflect the vague terms used in the text. With an important exception: we formalized “the relevant cloud service customer” via the predicate “PIIController”. The reason is that ISO 27018 includes a constitutive provision that defines the cloud service customer⁷.

Therefore, the cloud service customer *is* the PII controller, *i.e.*, the organization handling the data of the PII principals. In case the PII principals handle their data themselves, without any “broker” doing that on their behalf, they *are* the PII controller, *i.e.*, the data cloud service customer, that should receive the notification of unauthorized accesses from the PII processor.

The final ingredient needed to (4) and (2) are axioms relating to the predicates occurring in both, similar to the axiom in (1), which state that PII is by default considered as personal data:

$$\begin{aligned} \forall_x [(\text{PIIProcessor } x) \rightsquigarrow (\text{dataProcessor } x)], \quad \forall_x [(\text{PIIController } x) \rightsquigarrow (\text{dataController } x)], \\ \forall_x \forall_z [(\text{unauthorizedAccess } x \ z) \rightsquigarrow (\text{dataBreach } x \ z)], \\ \forall_x [(\text{promptly } x) \rightsquigarrow (\text{nonDelayed } x)] \end{aligned} \quad (5)$$

Axioms (5) are quite intuitive: for instance, the first one states that any entity that is taken to be a PII processor, with respect to the ISO 27018 standard, is also, by default, taken to be a data processor with respect to the GDPR.

⁷ISO 27018, Article 0.1: “The cloud service customer, who has the contractual relationship with the public cloud PII processor, can range from a natural person, a ‘PII principal’, processing his or her own PII in the cloud, to an organization, a ‘PII controller’, processing PII relating to many PII principals”.

It is easy to verify that every tuple of variables “ e ”, “ x ”, “ y ”, “ z ”, “ e_p ”, “ e_c ”, and “ e_b ” that satisfies formula (4) also satisfies formula (2) by default.

Again, such default correlations may be rewritten. For example, a judicial authority could later decide that although a notifying action may be considered as “prompt” with respect to the ISO 27018 standard, it cannot be considered as being “without undue delay” with respect to the GDPR. In such a case, the last implication in (5) needs to be overridden by higher-priority axioms: notifications would be henceforth considered “without undue delay” with respect to the GDPR only if they are “prompt” with respect to the ISO 27018 standard *and* they sport the extra features decided by the judicial authority.

Other instances of data breaches may be encompassed as well. For instance, ISO 27018 also imposes the prompt notification of “unauthorized access to processing equipment or facilities resulting in loss, disclosure or alteration of PII”. If we formalize such an unauthorized event “ e ” in terms of a predicate “(lossDisclosureOrAlteration $e z$)”, the corresponding formula is obtained by inserting this predicate in place of “(unauthorizedAccess $e_a z$)” in formula (4). The correlation with respect to the GDPR provision formalized in (2) is achieved by adding the following axiom to the TBox:

$$\forall_x \forall_z [(\text{lossDisclosureOrAlteration } x z) \rightsquigarrow (\text{dataBreach } x z)] \quad (6)$$

5 Discussion and Conclusion

In this paper we address the problem of legal compliance, *i.e.*, establishing how to be conforming to regulations and laws. The problem is not one of easy solution for at least two reasons. First, unless specific practices of compliance are directly mentioned in a law, what to do to achieve compliance is often unclear for an enterprise. Second, it is becoming hard to cope with a growing landscape of laws and regulations. To cope with such an overload, companies often resort to standards, which offer to the subject implementing them (not full compliance, but) a presumption of compliance. The question of establishing such a presumption thus reduces the problem to finding formal and substantive correlations between the terms used to describe the practices in the standards (*e.g.*, sending a digitally signed email to the PII controller) and those used in the provisions of the law (*e.g.*, duty of notification to the data controller).

This paper introduces a semi-automated methodology to help practitioners find and establish such correlations. The paper exemplifies the methodology in the context of data protection regulations (in particular, the GDPR) and security standards (ISO 27018), two domains which display significant overlaps. By following the illustrated methodology, one can build a machine-processable *knowledge base* of logic formulæ that model and store relevant concepts from a law and a standard, together with their possible formal correlations. The knowledge base, which will be collaboratively accessible, will have its records updated and labelled considering the outcomes of specific auditing processes or decision of the courts; in time, it will embed substantive correlations on which a legal argument for compliance can be relied upon. It should be pointed out that the logic into which we translate the correlations is defeasible, allowing them be overridden. The methodology herein is currently a work in progress and not fully implemented yet. The complete methodology requires the definition of a detailed taxonomy of concepts extracted from the law and the security standard. In the future, we envision significant developments of the research presented in this paper. In addition to the technical skills needed to manage the knowledge base, the methodology would greatly benefit from a close interaction with legal authorities. In a small country like Luxembourg, this could be easily achieved.

Several technical challenges related to building and updating the knowledge base are raised. The translations from natural language to logical formulæ must be uniform for excerpts of text that are similar to each other. To achieve this, we must overcome the limitations of a manual translation, which would be time-consuming and error-prone. For this reason, our work must rely on NLP technologies. However, even at the best of their performances, current NLP algorithms are still unable to automatically carry out the translation with a reasonable level of accuracy, so we advocate a *semi-automatic* translation of the provisions. Similar approaches are applied to translations in general, where translators are helped by collaborative tools such as the “SDL Trados Studio”⁸, which suggests, via pattern-recognition text-similarity NLP techniques [12, 3], how to translate a sentence on the basis of the translations of similar sentences that the translators have previously stored in the tool. Inspired by that approach, we will

⁸<http://www.translationzone.com/products/trados-studio/>

develop an enhanced text editor to assist the manual translation of provisions into formulæ. For each provision, the editor will display the translations of similar provisions found via NLP procedures applied to the provisions already stored in the knowledge base. Finally, we are aware that the knowledge base must be consistent, *i.e.*, without contradictions, even after having applied the defeasibility measures. To check for consistency, we plan to store formulæ using a XML-based data model, and employ/extend reasoners to monitor the consistency of the knowledge base, whenever new formulæ are added to it.

References

- [1] C. Bartolini, R. Muthuri, and C. Santos. Using Ontologies to Model Data Protection Requirements in Workflows. In *Proc. of the 9th Int. Work. on Juris-informatics (JURISIN)*, pages 27–40, Nov. 2015. Extended version to be published in LNAI book.
- [2] R. Benjamins, B. Selic, and A. Gangemi, editors. *Law and the Semantic Web: Legal Ontologies, Methodologies, Legal Information Retrieval, and Applications*, volume 3369 of LNAI. Springer-Verlag Berlin Heidelberg, 2005.
- [3] G. Boella, L. Di Caro, A. Ruggeri, and L. Robaldo. Learning from syntax generalizations for automatic semantic annotation. *The J. of Intelligent Information Systems*, 43(2):231–246, 2014.
- [4] D. Davidson. The logical form of action sentences. In N. Rescher, editor, *The Logic of Decision and Action*. Univ. of Pittsburgh Press, 1967.
- [5] P. De Hert and V. Papakonstantinou. The proposed data protection Regulation replacing Directive 95/46/ec: A sound system for the protection of individuals. *Computer Law & Security Review*, 28(2):130–142, April 2012.
- [6] P. De Hert, V. Papakonstantinou, and I. Kamara. The cloud computing standard ISO/IEC 27018 through the lens of the EU legislation on data protection. *Computer Law & Security Review*, 32(1):16–30, February 2016.
- [7] G. Governatori, F. Olivieri, A. Rotolo, and S. Scannapieco. Computing strong and weak permissions in defeasible logic. *Journal of Philosophical Logic*, 42(6):799–829, 2013.
- [8] J. Hansen. Prioritized conditional imperatives: problems and a new proposal. *Autonomous Agents and Multi-Agent Systems*, 17(1):11–35, 2008.
- [9] J. R. Hobbs. *Deep Lexical Semantics*, volume 4919 of LNCS, pages 183–193. Springer-Verlag Berlin Heidelberg, 2008.
- [10] J. Horty. *Reasons as Defaults*. Oxford University Press, 2012.
- [11] D. Makinson and L. W. N. van der Torre. Input/output logics. *Journal of Philosophical Logic*, 29(4):383–408, 2000.
- [12] R. Mihalcea, C. Corley, and C. Strapparava. Corpus-based and knowledge-based measures of text semantic similarity. In *Proc. of the 21st National Conference on Artificial Intelligence - Volume 1, AAAI’06*, pages 775–780. AAAI Press, 2006.
- [13] V. Reding. The upcoming data protection reform for the European Union. *International Data Privacy Law*, 1(1):3–5, February 2011.
- [14] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [15] L. Robaldo, L. Humphreys, L. Sun, L. Cupi, C. Santos, and R. Muthuri. Combining input/output logic and reification for representing real-world obligations. In *Post-proc. of the 9th Int. Work. on Juris-informatics. Lecture Notes in Artificial Intelligence*, 2016.
- [16] Karin Kipper Schuler. *Verbnet: A Broad-coverage, Comprehensive Verb Lexicon*. PhD thesis, Philadelphia, PA, USA, 2005. AAI3179808.