Early Studies in Acquiring Evidentiary, Reusable Business Process Models for Legal Compliance

Travis D. Breaux
Department of Computer Science
North Carolina State University
Raleigh, NC, USA
tdbreaux@ncsu.edu

Calvin Powers
IBM Software Group
Research Triangle Park, NC, USA
cspowers@us.ibm.com

Abstract

Government laws and regulations impose legal requirements on information practices in healthcare and finance. These requirements govern the use and disclosure of information across organizations and their business practices. To comply with the law, organizations must demonstrate that they have verifiable procedures in-place to implement these requirements. This paper surveys our experiences acquiring business process models expressed in the Business Process Model Notation (BPMN) using a systematic method. The method requires business process owners to classify regulatory statements using a legal ontology to identify legal requirements. The itemized requirements can then be used to specify elements in a business process model to demonstrate due diligence under the law. The contributions of this paper include lessons learned while acquiring the model with attention to traceability, distinguishing between legally expressed and implied activities and implementing legally imposed deadlines and suspensions. We discuss the lessons learned with examples from the U.S. Health Insurance Portability and Accountability Act (HIPAA).

1. Introduction

In the United States, government regulations require organizations to develop policies and procedures that comply with the law. These regulations are specified in complex and ambiguous legal language and generally prescribe business practices. For example, the U.S. Health Insurance Portability and Accountability Act[1] (HIPAA) Privacy Rule (42 C.F.R. §§160,164) was developed by the Department of Health and Human Services (HHS) to modernize health information systems. The HIPAA Privacy Rule affects some 580,000 different establishments in the US who employ over 14 million people [8]. Each of these establishments is responsible for interpreting the Privacy Rule and determining how their business practices comply with this law.

[1] U.S. Public Law No. 104-191 (1996)

The increasing reliance on software systems to support these practices presents business process owners, software engineers and system administrators with the daunting challenge of interpreting regulations to determine if their systems comply. In a 2007 Ersnt and Young survey of over 1,300 organizations, compliance with regulations was ranked as the top driver (64%) of information security and privacy [10]. The IT Policy Compliance Group, an association of IT vendors and audit associations benchmarked 876 organizations and found two activities clearly distinguish leaders from under-performers in compliance: (1) documenting assets, IT procedures and controls and (2) updating controls and procedures, such as in response to changing legal requirements [15]. Businesses need tools and methods to document and rationalize how their business processes align with the requirements of law.

In this paper, *legal compliance* means the ability to "maintain a defensible position in a court of law" [5]. Thus, compliance entails accumulating accurate and complete evidence that demonstrates *due diligence*, or "reasonable efforts that persons make to satisfy legal requirements or discharge their legal obligations" [11]. To support business process owners and software engineers in the exercise of due diligence, we are developing the Frame-based Requirements Analysis Method (FBRAM) for systematically acquiring legal requirements from policies and regulations [1]. The method has been previously applied to the HIPAA Privacy Rule [6, 2] and to the Section 508 Access Standards (36 C.F.R. §1194) [3].

In this paper, we examine how legal requirements acquired using the FBRAM can be used to create business process models. We believe that these models, together with traceability to and from the regulation text and codified legal requirements, provide auditors with a critical part of the reproducible and certifiable chain of evidence that shows how business processes and supporting software systems comply with laws.

The remainder of this paper is organized as follows: in Section 2, we review related work; in Section 3, we briefly introduce our legal ontology; in Section 4, we

show how to acquire legal requirements from regulations using our frame-based method; in Section 5, we describe how to create business process models from frame-based requirements; and we conclude with the discussion and summary in Section 6.

2. Related Work

Related work in requirements engineering to support legal compliance includes using formal models to perform specialized inference. Breaux et al. represent rights, permissions and obligations to balance rights with obligations [6] and to reason about goals using subsumption inference in Description Logic [4]. Massacci et al. extended Tropos, a goal modeling formalism, using DataLog to reason about acts of delegation and permission to check consistency in policy and law [19]. Miseldine et al. describe a model to support evidence-based compliance in business process outsourcing [20] and Karagiannis et al. illustrate how business process models can be used to support compliance with the Sarbanes Oxley Act [16]. Other models have been developed to manage traceability between legal texts and derived artifacts, a critical due diligence and software process requirement [17, 1, 12].

Methods to support compliance include a legal requirements acquisition method [6], which has associated tool-support to identify rights, permissions and obligations using manual [1] and automated annotation [18]. To support model checking of regulations, Delahaye et al. describe a method to identify hidden assumptions using the Focal environment [9]. Another challenge includes complex legal exceptions, which Breaux et al. address with a requirements prioritization method based upon priority hierarchies [2]. Finally, Breaux et al. have proposed requirements refinement patterns as a method to guide engineers in the refinement of legal requirements into product requirements [3].

Frameworks have also been proposed to manage traceability between goal models and legal documents [12] and manage accountability and traceability through refinement and delegation [7]. Finally, research in legal requirements has yielded reusable catalogues of privacy requirements [20] and information access constraints [2]. This paper extends this extensive body of work by investigating the relationships between legal requirements and business process models for the purpose of developing a systematic, repeatable method to align information system with relevant laws and maintain a chain of evidence to demonstrate due diligence.

3. Legal Requirements Ontology

The legal requirements upper ontology (see Figure 1) is used to classify statements in a legal text using statement-level concepts. These statement-level concepts are defined as follows:

- *Permission* (P) means any state that an entity is permitted to achieve, maintain or avoid, or any act that an entity is permitted to perform; permissions include stakeholder *rights*.
- *Obligation* (O) means any state that an entity is required to achieve or maintain, or any act that an entity is required to perform.
- *Refrainment* (R) means any state that an entity is required to avoid, or any act that an entity is required not to perform.
- *Exclusion* (E) means any state that an entity is not permitted or required to achieve, maintain or avoid, or any act that an entity is not permitted or required to perform.

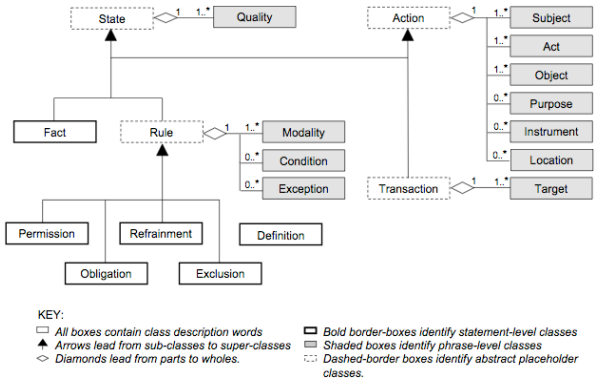

Figure 1: Legal requirements upper ontology

The FBRAM statement-level concepts correspond to the Hohfeld legal concepts for strict rights (or claims) and duties, which are similar to the concept of permission and obligation in our ontology, respectively [13]. Hohfeld identifies respective opposites that he calls no-rights and privileges, and which differ from how permissions and obligations are negated under the axioms of Deontic Logic [14]: "not permissible" implies "it ought not to be", sometimes called a prohibition or refrainment, which agrees with the Hohfeld concept of no-right. However, in Deontic Logic, "not obligatory" implies "is permissible", which Hohfeld calls a "privilege" and which is distinct and separate from a right or permission and is classified as an exclusion in the FBRAM. Whereas Deontic Logic can be used to reason about laws using a closed-world assumption, laws are written using an open-world assumption, evidenced by exclusions. Systems of legal inference must acknowledge this discrepancy.

In the next section, we consider how to apply these concepts to a sample legal text while maintaining traceability for auditing purposes.

4. Legal Requirements Acquisition

The Frame-based Requirements Analysis Method (FBRAM) was developed to provide a systematic process to acquire legal requirements from regulations while maintaining traceability [1]. The method employs the following four artifacts:

- the *upper ontology* containing legal concepts used to specify legal requirements, described in Section 3;
- *phrase heuristics* used to map legal statements and phrases to concepts in the upper ontology;
- a *context-free markup language* for annotating legal text with upper ontology concepts; and
- a *document model* for maintaining traceability between relevant paragraphs and phrases in legal text and the acquired legal requirements.

In this paper, we briefly describe the ontology, heuristics and markup language and refer the reader to our earlier work for a description of the complete method [1].

Laws and regulations like the HIPAA are written in dense legal language with numerous cross-references to other paragraphs, sections and other laws. In the Privacy Rule alone, there are over 400 cross-references. Figure 2 presents an excerpt from the HIPAA Privacy Rule §164.528(a)(1)-(2). The text is adapted precisely from the HIPAA Privacy Rule: elaborating text was omitted and markup (in bold) was added with no other changes made.

```
1  (1) {#P [#s The individual] [#m has a right
2      to] [#a receive] [#o an accounting of
3      disclosures made by the covered
4      entity], {#e except for [disclosures: [
5      (i)  To carry out treatment, payment and
6           health care operations as provided
7           in §164.502;&
8      (ii) To individuals of protected health
9           information about them as provided
10          in §164.502;& ... ]}}
11 (2)(i)   {#O [#s The covered entity] [#m
12          must] [#a temporarily suspend] [#o
13          an individual's right to receive an
14          accounting of disclosures to a
15          health oversight agency or law
16          enforcement official], {#c if [such
17          agency or official provides the
18          covered entity with a written
19          statement specifying the time for
20          which such a suspension is
21          required]}}.
```

Figure 2: Example legal text annotated with the frame-based markup language

The markup is used to map concepts from the upper ontology onto sentences and phrases in the legal text by encapsulating those phrases in curly "{ }" and square brackets "[ ]". Immediately after an open bracket, a concept code (a letter) from Table 1 follows a pound sign "#" and indicates that this concept is assigned to the text encapsulated in these brackets. For example, the markup "#P" on line 1 follows the opening curly bracket "{" and indicates a permission begins on line 1 and ends on line 10 where the closing curly bracket "}" appears. Phrases within a sentence are mapped to slots nested in a frame. For example, in Figure 2 the phrase "The individual" on line 1 is mapped to the subject slot in a permission frame (e.g., #s indicates a subject) to indicate that the individual is the actor who performs the action in this permission.

Table 1: Upper ontology concept codes used in the example legal text

Code	Concept	Code	Concept
P	Permission	a	Act
O	Obligation	o	Object
s	Subject	c	Condition
m	Modality	e	Exception

Existing tool support for the FBRAM includes a parser for the markup language that detects ambiguities including logical ambiguities, indicated by English conjunctions (and, or), and under-specification or missing slot assignments (e.g., every complete legal requirement must have a subject, action and object slot assignment). Tool support enables transforming frame-based requirements expressed in XML into an HTML template. Figure 3 presents the first legal requirement (a permission) acquired from §164.528(a)(1) formatted using the HTML template. The template presents the modality, natural language pattern extracted from the legal text, traceability information back to the legal text, and the assigned slot values. Analysts can use the markup to perform case splitting, as denoted by the dotted line in the exception slot in Figure 3.

Modality: Permission	
Pattern: [*subject*] [modality] [*act*] [*object*] {except for [*exception*]}	
Trace: ID: P1, Line 1:0, Source: 164.528(a)(1)	
Slots	**Values**
subject	individual
modality	has a right to
act	receive
object	an accounting of disclosures made by the covered entity
exception	*except for disclosures:* To carry out treatment, payment and health care operations as provided in §164.502. *except for disclosures:* To individuals of protected health information about them as provided in §164.502.

Figure 3: Example frame-based requirement created by parsing the annotated text

In Section 5, we describe how to map these slot assignments to elements in a business process model.

5. Codifying Business Processes

Business processes describe the tasks performed by stakeholders to fulfill a business need. Emerging business process modeling languages, such as the Business Process Execution Language for Web Services (BPEL4WS) and the Semantics for Business Vocabulary and Business Rules (SVBR) allow engineers to create machine-readable descriptions of web services and business processes. The Business Process Modeling Notation (BPMN) provides a visual interface to a subset of BPEL4WS, which can enable auditors, lawyers and engineers to evaluate how a company's business processes align with their legal requirements and service-oriented architecture. At the time of this writing, there were 53 publicly available, industry implementations of BPMN.

The remainder of this section describes the steps to specify a business process model in the BPMN using frame-based requirements. The examples that appear in this section were developed using 5 permissions and 14 obligations acquired by applying the FBRAM to §164.528 in the Privacy Rule describing the right of an individual to receive an accounting of disclosures.

5.1. Deriving activities from requirements

After the analyst has acquired frame-based legal requirements (see Section 4), they proceed to map each requirement to a set of business process elements in a BPMN diagram. Figure 4 illustrates a business process diagram expressed in BPMN that was acquired from permission P_1 shown in Figure 3. In BPMN, *activities* are represented by round-cornered rectangles and describe a unit of work. Activities are connected by *flows*, which describe the movement of information, called *message flows* and represented by arrows with dotted lines, and the movement of time, called *sequence flows* and represented by arrows with solid lines. Gateways, represented by diamonds, connect and control sequence flows using conditionally branch logic. The activities associated with a single participant (e.g., an individual) are grouped into a *pool* represented by a rectangular container spanning the length of the diagram. Pools are divided into one or more *swimlanes* that proceed chronologically from left to right and contain activities performed by the associated participant. Sequence flows pass between swimlanes, but only message flows pass between pools. Figure 4 shows only one swimlane for each participant.

The process to map the requirement from Figure 3 to the pools shown in Figure 4 proceeds in three steps: (1) identify the actor (the individual) in the subject slot of the frame-based requirement and create a pool and swimlane for this actor, if one does not already exist; (2) create an activity, represented by a round-cornered rectangle, by appending the object slot value "an accounting of disclosures…" to the act slot value "RECEIVE" – the act is capitalized for emphasis and constraints on the object may be omitted for simplicity; (3) create incoming gateways or activities for condition and exception slot values; and (4) annotate this derived activity with the legal requirement ID "P_1" using a comment represented by the open square bracket that connects to the activity via a dotted line.

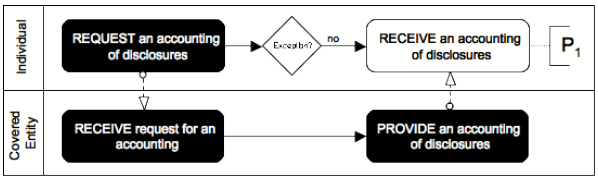

Figure 4: Example pools with legally expressed and implied activities.

5.2. Inferring activities from requirements

We distinguish between *expressed* activities systematically derived from legal requirements using the FBRAM and *implied* activities inferred from these requirements by the analyst. In this paper, we use our own convention of coloring expressed and implied activities with white and black backgrounds, respectively. For example, the expressed activity "RECEIVE an accounting of disclosures" was directly derived from permission P_1. This permission implies an obligation on another actor to "PROVIDE an accounting of disclosures". In Figure 4, the analyst inferred that this other actor was the "Covered Entity" and created a separate pool to contain the implied activity. This type of inference is called "balancing rights with obligations" and is systematically inferable from the frame-based requirement, given a list of transaction verb pairs (e.g., receive/ provide, give/ take, etc.) [6]. Because the implied activities are not expressed in the law, they must be checked with an appropriate auditor for review to determine if they are accurate legal interpretations.

5.3. Deriving deadlines from requirements

Legal requirements impose deadlines on performing certain obligations to improve accountability. Consider obligations O_4, O_7 and O_{11}, below, which require a covered entity to: suspend, provide and act on a request for provisions of an accounting of disclosures.

O_2: The covered entity *must* temporarily suspend an individual's right to receive an accounting of disclosures to a health oversight agency or law enforcement official. **§164.528(a)(2)(i)**

O_7: The covered entity *must* provide individuals with a written accounting. **§164.528(b)**

O_{11}: The covered entity *must* act on the individual's request for an accounting, no later than 60 days after receipt of such a request. **§164.528(c)(1)**

Figure 5 expands Figure 4 by changing the implied activity "PROVIDE an accounting of disclosures" to the expressed activity derived from obligation O_7. Obligations O_2 and O_{11} affect the sequence flow of other activities through suspension and deadlines, respectively; thus, special care is taken to implement these requirements in the BPMN. The BPMN provides notation for associating deadlines with activities using special events, called *timers*. The timer attached to the activity derived from O_7 implements obligation O_{11} by checking that this activity completes within 60 days. Any outgoing arrows from this timer would lead to an exception state that, when triggered, means the business process failed to comply with this obligation.

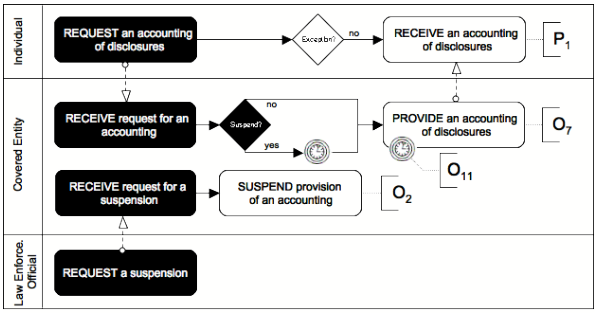

Figure 5: Example BPMN with deadlines on performing legal obligations

The BPMN version 1.1 specification does not support a "suspension" event. Therefore, the expressed activity "SUSPEND provision of an accounting" is derived from obligation O_2 and a gateway, represented by the diamond labeled "Suspend?" is introduced. This gateway is triggered, if "yes" a request for suspension has been received and acted upon, which delays the sequence flow to the provision of an accounting; otherwise, "no" a request was not received and the flow proceeds to the provision of accounting. The length of the delay is set to be the time specified in the suspension request received from law enforcement.

Contradictions in the law can be subtle. For example, obligation O_{11} requires covered entities to provide the accounting of disclosures to the individual no later than 60 days "after receipt of such a request", unless a one-time, 30-day extension is requested by the covered entity (not shown). If a law enforcement official requests a suspension for a period of time greater than 90 days, the covered entity will violate this deadline. Figure 6 illustrates this alternate interpretation, wherein the activity "PROVIDE an accounting of disclosures" is a composite process that contains inferred sub-processes or tasks. This time, the gateway implements obligation O_2 by delaying the sending of the accounting while obligation O_{11} continues to impose the deadline on the entire composite process. Thus, interpretation represented by this model could lead to violations of the law.

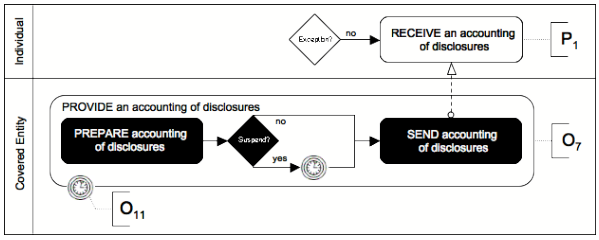

Figure 6: Example BPMN with conflicting deadlines on performing legal obligations

Executable languages that can check business process models for contradictions may catch this kind of vulnerability. However, the only recourse to resolving this contradiction is to study the Privacy Rule for relevant exceptions or to seek additional guidance from the Office of Civil Rights within HHS, which is responsible for HIPAA enforcement. Such ambiguity wedges business process owners between regulators and law enforcement, forcing them into the difficult position of making critical assumptions to "play it safe" by avoiding outcomes that expose their companies to potential legal violations.

6. Discussion and Summary

Increasingly, organizations must be able to demonstrate that they have verifiable procedures in-place to implement legal requirements imposed on information collection and use by government regulations and policies. Using examples from our in-depth analysis of the HIPAA, we have shown that business process models can be directly acquired from legal requirements. However, as previously discussed, analysts must still infer implied activities within the context of their business practices and resolve legal lacunae or ambiguous gaps in the law. To improve reliability in the acquisition process, business process owners need new methods to manage how they interpret laws in a controlled and systematic fashion. These methods have the potential to reduce non-compliance risk and increase assurance that systems comply with relevant policies and regulations.

Future work is needed to understand how to best visualize the implementation of legal requirements in governed lines of business. Business practices are created to fulfill a business need and are expressed in domain-specific language not always consistent with the language of governing laws. For example, an insurance company processes patient insurance claims, which is one line of business governed by the HIPAA Privacy Rule in the legal terms of covered entities and their uses and disclosures of protected health information. For example, are the activities in Figure 4

one-to-one and onto for an insurance company's business processes, or do their processes combine and/or divide these activities in complex ways? Further study is needed to know whether new BPMN extensions or design patterns are needed to support this potentially complex alignment.

We envision that some industries will benefit from business process models that are standardized from legal requirements and universally applicable, though with some re-engineering, to new and legacy information systems. To this end, the OMG Business Process Definition Metamodel (BPDM) may provide this capability by enabling analysts to reuse industry-standard business process classes. The extent to which this vision requires specialized architectures, such as those currently being investigated in the web service community, is yet to be seen. Because U.S. regulators are reluctant to impose harsh and overburdening restrictions on businesses for fear of stifling innovation, this vision will likely need to be studied in an industrial or academic setting.

7. Acknowledgements

This work was supported in part by the IBM PhD Fellowship (RTP-CAS).

8. References

[1] T.D. Breaux, A.I. Antón, "A systematic method for acquiring regulatory requirements: a frame-based approach", *6th International Workshop on Requirements for High Assurance Systems (RHAS-6), Delhi, India,* Sep. 2007.

[2] T.D. Breaux, A.I. Antón, "Analyzing regulatory rules for privacy and security requirements", *IEEE Trans. Soft. Engr.*, 34(1): 5-20, Jan./Feb. 2008.

[3] T.D. Breaux, A.I. Antón, K. Boucher, M. Dorfman, "Legal requirements, compliance and practice: an industry case study in accessibility", *16th IEEE Int'l Conf. on Req'ts Engr.*, pp. 43-52, 2008.

[4] T.D. Breaux, A.I. Anton, J. Doyle, "Semantic parameterization: a conceptual modeling process for domain descriptions," *ACM Trans. Soft. Engr. Methods.*, 18(2): 5, Nov. 2008.

[5] T.D. Breaux, A.I. Antón, C.-M. Karat, J. Karat, "Enforceability vs. accountability in electronic policies", *IEEE 7th International Workshop on Policies for Distributed Systems and Networks (POLICY'06)*, London, Ontario, pp. 227-230, Jun. 2006.

[6] T.D. Breaux, M.W. Vail, A.I. Antón, "Towards compliance: extracting rights and obligations to align requirements with regulations," *14th IEEE Int'l Conf. on Req'ts. Engr.*, pp. 49-58, 2006.

[7] T.D. Breaux, A.I. Antón, E.H. Spafford, "A distributed requirements management framework for compliance and accountability", (In Press) Computers and Security, Oct. 2008.

[8] Bureau of Labor Statistics, U.S. Dept. of Labor, "Health Care," Career Guide to Industries, 2008-2009.

[9] D. Delahaye, J.-F. Etienne, V.V. Donzeau-Gouge, "Reasoning about airport security regulations using the focal environment", *2nd Int'l Symp. Leveraging Applications of Formal Methods, Verification and Validation*, Paphos, pp. 45-52, Nov. 2006.

[10] Ernst and Young, 10th Annual Global Information Security Survey: Achieving a Balance of Risk and Performance", 2007.

[11] B.A. Garner, editor. Blacks Law Dictionary, 8th ed. ThompsonWest, St. Paul, Minnesota, 2004.

[12] S. Ghanavati, D. Amyot, L. Peyton, "Torwards a framework for tracking legal compliance in healthcare," *19th Int'l Conf. Adv. Info. Sys. Engr.*, pp. 218-232, 2007.

[13] W.N. Hohfeld. "Some fundamental legal conceptions as applied in judicial reasoning." *The Yale Law Journal*, 23(1):16–59, 1913.

[14] J.F. Horty, *Agency and Deontic Logic*, Oxford University Press, New York NY, 2001.

[15] IT Policy Compliance Group, "Managing Spend to Improve Compliance Results", Oct. 2006.

[16] D. Karagiannis, J. Mylopoulos, M. Schwab, "Business process-based regulation compliance: the case of the Sarbanes-Oxley Act", *15th Int'l Req'ts Engr. Conf.*, Delhi, India, pp. 315-321, 2007.

[17] S.L. Kerrigan and K.H. Law. "Logic-based regulation compliance-assistance." *9th Intl Conference on Artificial Intelligence and Law*, pp. 126–135, Edinburgh, Scotland, 2003.

[18] N. Kiyavitskaya, N. Zeni, T.D. Breaux, A.I. Antón, J.R. Cordy, L. Mich, J. Mylopoulos, "Automating the extraction of rights and obligations for regulatory compliance", *27th Int'l Conf. Conceptual Modelling (ER'08)*, Barcelona, Spain, pp. 154-168, Oct. 2008

[19] F. Massacci , J. Mylopoulos , N. Zannone, "Computer-aided support for Secure Tropos," *Automated Software Engineering*, 14(3): 341-364, Sep. 2007.

[20] P.L. Miseldine , U. Flegely, A. Schaad, "Supporting evidence-based compliance evaluation for partial business process outsourcing scenarios", (In Press) *1st Int'l Workshop on Req'ts Engr. and Law*, Barcelona, Oct. 2008.

[20] S. Toval, A. Olmos, M. Piattini. "Legal requirements reuse: a critical success factor for requirements quality and personal data protection," *IEEE Int'l Conf. Req'ts Engr.*, pp. 95-103, 2002.