Analyzing Policy Commitments for Privacy and Security Requirements

Jessica D. Young and Annie I. Antón

North Carolina State University
College of Engineering
Department of Computer Science
Raleigh, NC, USA

{jdyoung2, aianton}@ncsu.edu

http://www4.ncsu.edu/~{jdyoung2, aianton}

Abstract Companies that publish privacy policies espouse to protect customers' sensitive information. These privacy policy documents state company-specific commitments to the customer about specific data and how the company will collect, use and securely store it. It is incumbent upon requirements engineers to understand these commitments so that they can be operationalized into specific security and privacy requirements. It is imperative that companies ensure that these commitments are maintained in the software systems that are governed by corporate policies. This paper proposes the analysis of policy commitments for privacy and security requirements within the context of U.S. health care institutions' websites.

Keywords: privacy requirements, security requirements, policy, commitments, rights, legal obligations

1 Introduction

Identifying privacy and security requirements for systems that are governed by policy and law is increasingly important as failure to comply with law can result in stiff penalties. This paper examines the efficacy of applying Semantic Parameterization [BVA06] to extract rights and obligations from health care institutions' website privacy policy documents. Prior work established that privacy policies are useful sources for extracting security and privacy requirements for a system [AEC03]. Breaux et al. extracted rights and obligations to analyze law and extract requirements [BVA06]. Otto and Antón explained the importance of having traceability from law to policy to software requirements in order to establish a basis for regulatory compliance [OA07].

Herein, we investigate whether the rights and obligations that are expressed in law could be tied to the rights and obligations expressed in policy documents. Our analysis reveals that a majority of statements made in privacy policy documents actually express commitments from the organization to the customer as well as

customer and organizational rights. Our experience suggests that rights and obligations are a more appropriate unit for analyzing the law for software requirements than for analyzing privacy policy documents for software requirements. Our findings suggest that privacy policy documents predominantly express commitments and rights, rather than legal obligations. Using grounded theory [GS67], we have developed a theory of commitments for identifying software requirements.

In the United States, the Federal Trade Commission (FTC) protects consumers from unfair business practices through the Federal Trade Commission Act [FTC14]. This act allows the FTC to enforce the privacy policies that organizations post on their websites with regard to how they collect, keep, use, and preserve their customers' personal information [Hof06]. Therefore, the practices an organization commits to in its policies become required and expected. When we are extracting commitments, rights, and legal obligations from the policy documents we are looking at what the organization was required to put in the policy (legal obligations) versus what they chose to put in the policy (commitment).

A *commitment* is a pledge that an organization makes to its customers. At the time a commitment is made law does not require it. Once a commitment is articulated and made public in the form of a policy the U.S. FTC expects the organization to abide by the commitment. If a commitment is violated then the FTC views that as an "unfair and deceptive" practice that is subject to sanction [FTC14]. A *right*, organizational or customer, is an action that the stakeholder is entitled to perform. A *legal obligation* is an action that an organization is legally bound to perform. Jackson describes a requirement as being "a condition over phenomena of the environment" [Jac97]. Commitments can be operationalized as software requirements.

We take as our starting point, the intuitive and highly loaded concept of a "right," because privacy management tools are rights management tools. Policies grant various rights to different parties, such as the right to refrain from disclosing information, the right to use information freely given and consented to, the right to know whether information is being gathered, etc. Much has been written about the nature of rights in general. For example, whether a constitutional right to privacy exists in the United States has been a perennial subject of controversy, the

current legal position in the U.S. being that "privacy" means freedom from undue interference (Judge Brandeis's famous "right to be let alone" [WB90]).

Our position is that rights management is only feasible if one stipulates an operational definition of what a right is. Accordingly, we adopt the perspective articulated by O'Neill [ONe02] in her discussion of autonomy rights and trust in bioethics: "A right is a relational claim legitimately ascribed to a right bearer with respect to an implicit or explicit other, the counterparty. To say that a party has a right is always a way of talking about the counterparty's implied obligation" [BVA06]. Thus, if I have a right to breathe unpolluted air, this means (among other things) that you are prohibited from smoking in my presence; if I have a right to privacy, it means that you are not allowed to intrude on me. Reformulating rights in terms of obligations of counterparties has important consequences for rights management: (1) the vague and contentious language of rights is made concrete and applicable to descriptions of states and actions, the matter of system specifications; (2) instead of a party "having" a right, we are forced to think about relational commitments and communication among parties, the very matter of information systems; (3) by defining rights and commitments as relationships between parties, we can specify and analyze not only rights and commitments themselves but their conditions of legitimacy, questionability and relativity to circumstances.

The remainder of this paper is organized as follows. Section 2 provides an overview of relevant work. Section 3 discusses the methodology for this study. Section 4 includes results and discussion of our study. Section 5 describes future work.

2 Related Work

Our work builds upon several areas of research: requirements engineering, commitment theory, intentions, and contracts. Section 2.1 discusses privacy policy documents as a source of software requirements and the need to maintain traceability across law, policy, and software requirements in order to demonstrate due diligence in developing policy-compliant systems. Previous work related to commitment theory is discussed in Section 2.2, while Section 2.3 discusses intentions and contracts. The related work is important to understand in order to properly analyze policy documents because rights and obligations are only

suitable for describing a subset of the kinds of statements expressed in privacy policy documents.

2.1 Requirements Engineering

Policies and software requirements are similar in that they both express desire or worth [AEC03]. Policies, however, are broader in scope than software requirements because policies can govern multiple systems and software requirements are typically specified for one system [AEC03]. In addition, policies are more open-ended [AEC03]. This work provides support of our work and shows that policies can be useful during requirements phase.

Robinson states that there is a need for software requirements to comply with policies [Rob05]. Based on this need, he developed a framework, REQMON, to monitor software requirements at runtime [Rob05]. Ghanavati et al. explore ways to develop legally complaint systems; they examined the use of three types of links—traceability links, compliance links, and responsibility links [GAP07a, GAP07b]. In particular, software that is covered by law must comply with all governing laws and policies. In order to be able to establish such compliance, traceability from relevant regulations to requirements specifications is essential [BA08, OA07]. In the event of a security breach, such traceability enables requirements engineers to ensure that all software artifacts are auditable and usable in a court of law to demonstrate due diligence. Our work supports traceability based on commitments, rights, and legal obligations extracted from policy documents.

Recently, Breaux and Antón introduced KTL, a context-free grammar they used to analyze the most frequently expressed goals in over 100 Internet privacy policy documents [BA05]. Semantic parameterization, a process for representing domain descriptions in first-order predicate logic [BAD09] was applied to goals and law but not to the original natural language policy documents [BA05]. As such, we attempt to apply this approach to natural language privacy policy documents rather than using an intermediate representation of goals as has been done previously. In addition, the ability to balance rights and obligations in regulations is critical for software engineers because a right that is not balanced by an obligation cannot be operationalized as a software requirement [BVA06]. Our work investigates whether this balance is also present in policy documents.

2.2 Commitment Theory

A commitment is essentially a pledge to do something. Within the context of privacy policies, a commitment is pledge to collect certain kinds of information for a specific purpose, use that information accordingly, and to store sensitive information for a certain length of time and with specific protections.

Several researchers have studied commitments from various points of view as we now discuss. Haddadi examines commitments as they relate to agents; he discusses four types of conditions: maintenance conditions, formation conditions, revision conditions, and responsibilities [Had95]. Maintenance conditions are conditions that must be maintained throughout the commitment [Had95]. Formation conditions are the conditions under which the commitment is formed [Had95]. Revision conditions are conditions under which an agent should "re-consider his commitment" [Had95]. Responsibilities are what the agent should do "after he gives up a commitment" [Had95]. Although his work primarily focuses on agents, his formation conditions relate to our work; our formation conditions include legal, transactional, data, and contractual conditions as expressed in privacy policy documents.

Wan and Singh examine multiagent systems and the use of commitments within multiparty agreements; they propose the following definition, "A commitment is an obligation from a debtor *x* to a creditor *y* about a particular condition *p*" [WS05]. They state that commitments have two forms—unconditional and conditional [WS05]. We adopt this distinction between conditional and unconditional commitments in our analysis of privacy policy documents for software requirements. As we discuss in Section 3, policy statements that express a commitment with a condition are indicators of a conditional commitment. For example, the terms "if," "unless," or "except" suggest a conditional commitment. Consider this statement from Aetna's Web Privacy Statement[1]: "When you visit and navigate our sites and when you communicate with us via our sites, we will not collect personal information about you unless you provide us that information voluntarily." The condition for the commitment that corresponds to this statement would be "unless you provide us that information voluntarily."

From a marketing perspective, Garbarino and Johnson consider overall satisfaction, trust, commitment, and future intentions for transactional and

relational marketing; they do so by surveying customers about each for an off-Broadway theater company in New York [GJ99]. They define commitment, based on the definition from Moorman, Zaltman, and Deshpandé, as "an enduring desire to maintain a valued relationship" [GJ99, MZD92]. Garbarino and Johnson asked the customers whether they strongly disagree to strongly agree to the following statements using a Likert scale: "I am proud to belong to this theater," "I feel a sense of belonging to this theater," "I can care about the long-term success of this theater," and "I am a loyal patron of this theater" [GJ99]. Their objective is to better target marketing efforts to customers to ensure they continue to be theater patrons in the future. In contrast, we are interested in examining commitments to ensure that software fulfills the commitments to the customer.

Several researchers have also examined commitments within the context of employee satisfaction and customer perception of employees. Baugh and Roberts investigated the job performance and commitment of engineers to determine whether engineers are more committed to their profession or to their organization [BR94]. Their studies reveal that rather than being contradictory or incompatible the connection between the two types of commitment—organizational and professional—may be interdependent [BR94]. Similarly, Reid et al. examined individuals' desire to continue to be employed by a government agency by looking at the employees' commitment to their work and the specific government agency [RRA06]. Ning-jun et al. examined how a service employee's commitment to the organization relates to customer perception, they believe that a fulfilled employee yields a happy customer, but further theoretical proof and empirical research is needed [NYQ07]. This relates to our work because a customer's satisfaction can be influenced by an organization's commitments within its policy.

Researchers have examined commitments within the context of duties or responsibilities that have to be filled. Tarhan et al. consider commitments within the software development process because the tasks of software engineers are often divided in terms of the commitments to complete the tasks that are assigned different developers [TDD99]. They created the Distributed Commitment Management Tool (DCMT) to help developers keep track of their commitments within the software development process, including goals, tasks, responsibilities,

[1] http://www.aetna.com/about/privacy.html, accessed August 21, 2008.

and resources [TDD99]. Similarly, Maggs discusses an individual's commitment to completing a task [Mag99]. Our work is similar because we are examining the commitments that organizations make to their customers as well as the responsibilities that organizations assume to fulfill those commitments.

2.3 Intentions and Contracts

An intention is a future plan or event. A contract is an agreement between two parties. In the United States, the Federal Trade Commission Act protects consumers from "unfair or deceptive acts or practices" [FTC14]. The FTC enforces the privacy policies that organizations post on their websites about how they collect, keep, use, and preserve their customer's personal information [Hof06]. Within the context of enforcing these fair practices within commerce, Beales discusses the history of the unfairness authority that enacts laws to protect against the invasion of 'public policy' [Bea03].

The 1987 *Intentions and Plans in Communication and Disclosure* workshop was comprised of a vast range of research fields, including computer science, artificial intelligence, and linguistics; this group of researchers contributed to the theories of communication as related to agents [CMP90]. In particular, Cohen and Levesque noted the need to state the required 'rational balance' between the beliefs, goals, plans, intentions, commitments, and actions of autonomous agents [CL90a, CL90b]. Cohen and Levesque presented a tiered formalism with the following operators: BELief, GOAL, HAPPENS, and DONE [CL90b]. Most relevant to our work are the persistent goals, which relate to commitments and intentions. According to Cohen and Levesque, an agent should believe that it can achieve all of its intentions [CL90a, CL90b]. Levesque et al. extend this work by examining the commitment of a group to a common goal, which they call a joint commitment [LCN90]. This relates to statements within policy documents that state that multiple organizations or portions of organizations are committed to doing a task.

Singh notes some limitations with Cohen and Levesque's work [CL90a], including "counterintuitive results" and "conceptual shortcomings" [Sin92]. Singh mentions a particular assumption ("all goals are eventually dropped") in combination with definition, which has an outcome that is counterintuitive [Sin92]. Shortcomings are present based on claims that were not proved [Sin92]. Singh and Asher explain that designers need a theory of intentions to describe the

behavior that is expected of the agents [SA91]. Formal language and semantics for intentions include: satisfaction conditions, content, honesty, and subsumption conditions [SA91]. Future events are characterized by these intentions [SA91]. Wu "demonstrated through tracing detailed examples that an integrated model of discourse does encompass discourse structure and intentions" [Wu93]. Privacy policies contain the intentions of the organizations.

Macneil gave eleven factors relating to contracts, included in this list was obligations [Mac80]. He explains that obligations vary in three aspects: *sources of content*, *sources of obligation*, and *specificity* [Mac80]. This relates to our work in that each obligation contains different content based on what is required and different sources based on which law requires the obligation.

3 Methodology

This section discusses our methodology for this study. Section 3.1 discusses the materials used—the policy documents. Section 3.2 discusses the execution of our methodology that includes extracting statements from privacy policy documents, classifying the identified statements as commitments, rights, and legal obligations, and documenting each statement using a template. We decompose each document into individual statements, which are typically a single sentence but some statements contains multiple sentences as we discuss in Section 3.2.1.

3.1 Materials

The companies whose policy documents we are examining are the same companies that were analyzed by Antón et al. in their work that mined goals from privacy documents of health care institutions [AEV07]. In this paper, we discuss policies of Aetna[2], a health care benefits company, and drugstore.com[3], an online drugstore. The policy documents—privacy policies and notices—are listed in Table 1. These documents ranged from one page to fourteen pages in length. While Aetna has many documents, the Notice of Privacy Practices documents were extremely similar for the seven different health care benefits plans listed in Table 1—Employee Assistance Plan, Flexible Spending Account Debit Card, Insured Health Benefits Plan, Long-Term Care Plan, Rx Home Delivery, Strategic Resource Company, and Student Health Plan. For the most part these notices had

[2] http://www.aetna.com
[3] http://www.drugstore.com

many of the same commitments, rights, and legal obligations and only differed based on the plan name, for example, Employee Assistance Plan or Student Health Plan. Aetna has thirty pages of policy documents while drugstore.com has twenty-seven pages worth of policy documents. The reason the Pharmacy Notice of Privacy Practices for drugstore.com is so many pages long is that it includes specific obligations for each state based on the state's laws.

Table 1: Aetna and drugstore.com Policy Documents

Aetna	drugstore.com
Notice of Privacy Practices by Plan Type • Employee Assistance Plan[4]—three pages • Flexible Spending Account (FSA) Debit Card[5]—four pages • Insured Health Benefits Plan[6]—four pages • Long-Term Care Plan[7]—four pages • Rx Home Delivery[8] (mail order pharmacy)—two pages • Strategic Resource Company (SRC)[9]—three pages • Student Health Plan[10]—four pages Notice of Information Practices by Plan Type • Large Case Pension[11]—one page • Life and Disability[12]—two pages • Privacy Notices[13]—one page • Web Privacy Statement—two pages	• Pharmacy Notice of Privacy Practices[14]—fourteen pages • Privacy Policy[15]—six pages • Terms of Use[16]—seven pages

3.2 Execution of Methodology

Our process initially entailed extracting rights (as defined by Breaux et al. [BA05, BVA06]) and legal obligations (as defined herein) from the policy documents listed in Table 1. It quickly became apparent that each document contained statements that actually emphasize the commitments that organizations make to their customers rather than pure rights and obligations as found in

[4] http://www.aetna.com/data/68164.pdf, accessed August 21, 2008.
[5] http://www.aetna.com/data/68306.pdf, accessed August 21, 2008.
[6] http://www.aetna.com/data/67806.pdf, accessed August 21, 2008.
[7] http://www.aetna.com/data/67969.pdf, accessed August 21, 2008.
[8] http://www.aetna.com/about/pdf/68059.pdf, accessed August 21, 2008.
[9] http://www.aetna.com/data/68113.pdf, accessed August 21, 2008.
[10] http://www.aetna.com/about/pdf/68050.pdf, accessed August 21, 2008.
[11] http://www.aetna.com/about/pdf/63510.pdf, accessed August 21, 2008.
[12] http://www.aetna.com/about/pdf/68049w.pdf, accessed August 21, 2008.
[13] http://www.aetna.com/about/information_practices.html, accessed August 21, 2008.
[14] http://www.drugstore.com/npp, accessed August 21, 2008.
[15] http://drugstore.custhelp.com/cgi-bin/drugstore.cfg/php/enduser/std_adp.php?p_faqid=191, accessed August 21, 2008.
[16] http://drugstore.custhelp.com/cgi-bin/drugstore.cfg/php/enduser/std_adp.php?p_faqid=512, accessed August 21, 2008.

regulations. For example, many statements begin with the words "we will." Within a policy Campbell states, "The word *will* means that [the organization is] committed to that position or action" [Cam98]. This suggests a commitment rather than a right or an obligation. Using grounded theory, we observed the following types of elements that form the basis for using commitments to reason about requirements that satisfy statements made in policy documents. These elements are defined as follows:

- A *commitment* is a pledge that an organization makes to its customers. At the time a commitment is made, it is not required by law. Once a commitment is articulated and made public in the form of a published policy the U.S. FTC expects the organization to abide by the commitments it made in the policy. If a commitment is violated then the FTC views that as "unfair and deceptive acts or practices" that is subject to sanction [FTC14].
- A *legal obligation* is an action that an organization is legally bound to perform.
- An *organizational right* is an action that an organization is entitled to perform.
- A *customer right* is an action that a customer is entitled to perform.

For the extraction process, we annotated each commitment, right, or legal obligation using the following template. The template was maintained in tabular form, using a Microsoft Excel workbook.

- an *ID* or identification code that we assigned to the statement,
- the *policy statement* as expressed in the policy document,
- the *policy document* is the name of the policy document that the statement is from,
- the *type* for the statement—commitment, right, or legal obligation,
- the statement's *responsible stakeholder* or subject (e.g., organization, member of organization),
- the statement's *action* (e.g., use, disclose, obtain, protect),
- the *object* of the statement's action (e.g., personal information),
- the *source* of the statement's object or "the provider" [BVA06] (e.g., organization, member of organization),

- the *target* of the statement's action or as Breaux et al. described it "the receiver" [BVA06] (e.g., organization, member of organization),
- the *purpose* for the statement (e.g., comply with the Health Insurance Portability and Accountability Act),
- the *conditions* within the statement (e.g., if the request for amendment is denied), and
- the *scenarios* or examples within the statement.

Some of these attributes are adopted from the GBRAM [Ant96] and Breaux et al. [BVA06], including: responsible stakeholder, action, object, source, target, and purpose. Consider the following example statement from a policy document that expresses a right (how each item was extracted will be described in Section 3.2.3):

Example 1:
ID – 1.r
Policy statement –"We may use your PHI [Protected Health Information] to treat you. For example, if you are being treated for an injury, we may share your PHI with your primary physician so they can provide proper care."
Policy document – drugstore.com Pharmacy Notice of Privacy Practices
Type – organizational right
Responsible stakeholder – drugstore.com
Action – use
Object – drugstore.com member's PHI
Source –
Target –
Purpose – treat drugstore.com member
Conditions –
Scenarios –if you are being treated for an injury, we may share your PHI with your primary physician so they can provide proper care.
Right – use drugstore.com member's PHI to treat drugstore.com member

The execution of our methodology involves three main steps: (i) breaking the policy document into statements, (ii) classifying the identified statements as commitments, rights, and legal obligations, and (iii) documenting each statement using a template. In the following subsections we describe these three steps.

3.2.1 Extracting Statements from Policy Document

While most statements are easily extracted from the policy document by making each heading and sentence a separate statement, some statements actually contain multiple sentences. If the document contains lists or examples, then breaking the document into statements is not as simple as making each sentence into a separate statement. Breaux and Antón addressed a similar problem, called continuations, within regulations that contain statements that break across subparagraphs [BA07].

Within the policy documents there are three types of lists based on the presence of headings and complete sentences in the items. In a *Type 1* list each item does not have a heading. A *Type 2* list's items have headings followed by complete sentences, whereas a *Type 3* list has items with headings that are followed by an incomplete sentence.

Type 1 List: The items of the list do not have headings. If the items of the list do not have headings, then the statement before the list would be prepended to the first sentence if each list item. This was the best solution to the problem because the items of the list were always incomplete sentences and required the statement before the list in order to make sense.

Type 2 List: The items of the list have headings that are followed by complete sentences. If the items of the list have headings that are followed by complete sentences, then we include the item heading with the statement before the list. This solves the problem of having an incomplete sentence before the list.

Type 3 List: The items in the list have headings that are followed by incomplete sentences. If the items in the list have headings that are followed by incomplete sentences, it is a continuation of the statement before the list. This requires the requirements engineer to combine what is done for *Type 1* and *2* lists. The statement before the list is prepended to the incomplete sentence in each of the items in the list. After this, the headings of the items are included in the statement before the list. This solves the problem of having an incomplete sentence before the list and incomplete sentences within the list items. When breaking the document into statements, examples within the document require more than just dividing the document by sentences.

Examples: Whereas most policy statements comprise a single sentence, some statements contain multiple sentences. Typically, the additional sentences contain examples to supplement the statements. In this case, the sentences are combined into one statement because alone the examples are not a commitment, right, or legal obligation, but instead the examples are the scenarios for a commitment, right, or legal obligation contained in a different sentence. Consider this statement (two sentences) from Aetna SRC Notice of Privacy Practices: "We may disclose information to doctors, dentists, pharmacies, hospitals, and other health care providers who take care of you. For example, doctors may request medical

information from us to supplement their own records." The example describes one of the possible scenarios.

3.2.2 Classifying the Identified Statements as Commitments, Rights, and Legal Obligations

When examining a statement from a policy document requirements engineers need to ask several questions to classify the statements. We now present the questions that are raised for identifying the statement as a commitment, a right, or a legal obligation.

Commitments. To identify commitments, we examine privacy policy statements by posing the question, "*Does this statement contain a commitment from the organization to the customer?*" If the statement does contain a commitment, "*What is the commitment?*" or "*What is the responsible stakeholder committing to do?*"

We found that commitments were often formatted as "[responsible stakeholder] will [action]…" Consider the following statement from the drugstore.com Pharmacy Notice of Privacy Practices: "We will include a copy of both statements in your file." The commitment drugstore.com makes to its members is *include a copy of both statements in drugstore.com member's file*.

Although the word "will" often signals a commitment there is an exception if the action is actually required by law. These statements are classified as legal obligations. Consider this statement from the drugstore.com Pharmacy Notice of Privacy Practices: "We will disclose your PHI when required to do so by local, state or federal law, including workers' compensation laws." Although "will" was used the statement is actually a legal obligation because law requires the disclosure of drugstore.com members' PHI. The legal obligation is: *disclose member's PHI when required to do so by local, state or federal law, including workers' compensation laws*.

Rights. To identify rights, we examine statements by asking the question, "*Does this statement contain a right?*" If the statement does contain a right, "*What is the right?*", "*What does the responsible stakeholder have the right to do?*", or "*What may the responsible stakeholder do?"*

We found that rights were usually formatted in one of two ways: "[responsible stakeholder] may [action]…" or "[responsible stakeholder] has the right to [action]…" The right is also classified by who the responsible stakeholder is—the

organization or the customer. Consider this statement from drugstore.com Pharmacy Notice of Privacy Practices: "We may use your PHI to treat you." This statement expresses drugstore.com's right to *use drugstore.com member's PHI to treat the drugstore.com member*. Another example from the same policy document is: "You have the right to request that we restrict how your PHI is used or disclosed in carrying out treatment, payment, or health care operations." The customer right indicated in this statement is *request that drugstore.com restrict how drugstore.com member's PHI is used or disclosed*.

Legal obligations. To identify legal obligations, we examine statements by asking the question, "*Does this statement include a legal obligation?*" If the statement does contain a legal obligation, "*What is the legal obligation?*" or "*What is the responsible stakeholder required by law to do?*" Frequently, these statements cited the specific law that requires the obligation.

We found that legal obligations were usually formatted in one of the following ways: "[responsible stakeholder] must [action]…" or "[responsible stakeholder] is/are required to [action]…" Consider this statement from Aetna Privacy Notices: "Aetna is required to send a notice ("Notice of Privacy Practices") to members of our insured Health and Long Term Care plans and Mail Order Pharmacy customers." The legal obligation, as required by the Health Insurance Portability and Accountability Act (HIPAA), is *send a notice ("Notice of Privacy Practices") to Aetna's insured Health and Long Term Care plans members*.

As mentioned above, some legal obligations may appear to be commitments but actually say that law requires the action.

3.2.3 Documenting Each Statement Using a Template

We now discuss how the attributes listed in the beginning of Section 3.2 are extracted from a statement. When extracting these attributes, any pronouns should be changed to the antecedent that the pronoun is referencing. We replaced pronouns that were referring to the organization with the "[organizations name]" (e.g., Aetna). Those pronouns that refer to the member are replaced with "[organization name] member" (e.g., Aetna member).

The requirements engineer needs to identify *ID, policy statement, policy document, type, responsible stakeholder, action, object, source, target, purpose, conditions*, and *scenarios* from each given privacy policy statement. It is important to note that most statements will not contain every attribute—it is

important to identify all relevant attributes. We demonstrate this extraction process for Example 1 from the beginning of Section 3.2.

ID. An identification code is needed in order to easily talk about a particular commitment, right, or legal obligation. The identification code is assigned to the statement based on how it is classified. Identification codes for commitments contain a *c*, rights contain an *r*, and legal obligations contain an *o*. Along with the letter representing the type each statement is also assigned a unique number. Since the statement in Example 1 was identified as a right its ID contains an *r*. It was assigned the number one because it was the first example of the extraction process we gave. Therefore, the statement's ID is 1.r.

Policy statement. The policy statement is the statement as it is taken from the policy document using the rules described in Section 3.2.1. It is important to document the policy statement along with the commitment, right, or legal obligation in order to know how it was stated within the policy document and to maintain traceability as well as the relevant context for future analysis. The policy statement for Example 1 is "We may use your PHI to treat you. For example, if you are being treated for an injury, we may share your PHI with your primary physician so they can provide proper care."

Policy document. The policy document is simply the name of the policy document that the statement came from. The document name is needed in order to know where the commitment, right, or legal obligation came from and for traceability reasons. For Example 1, the policy document is the "drugstore.com Pharmacy Notice of Privacy Practices."

Type. The type for the statement is identified as a commitment, right, or legal obligation. The type is identified as described in Section 3.2.2. The statement in Example 1 answers the question, "*What may the responsible stakeholder do?"* and is in the format of "[responsible stakeholder] may [action]…" The responsible stakeholder is drugstore.com; therefore the statement is classified as an organizational right.

Responsible stakeholder. The responsible stakeholder in the statement is the subject of the statement, who is performing the action. As previously mentioned, if the subject of the statement is a pronoun then the antecedent for the pronoun must be listed as the responsible stakeholder. It is necessary to document the responsible stakeholder to capture who is making the commitment, has the right,

or has the legal obligation. The responsible stakeholder in Example 1 is "drugstore.com."

Action. The action of the statement is the action that the responsible stakeholder is performing. The action that drugstore.com is performing in Example 1 is "use." The action is important because it will be the action of the commitment, right, or legal obligation.

If a statement has multiple actions within it, each action in the statement should be classified according to whether it conveys a commitment, a right, or a legal obligation. This is necessary because each commitment, right, and legal obligation can only have one action associated with it. Consider the following statement from the drugstore.com Pharmacy Notice of Privacy Practices: "We will use and disclose PHI to carry out health care operations." This statement expresses two commitments, which we document as follows: (1) *use PHI to carry out health operations* and (2) *disclose PHI to carry out health care operations*.

Object. The object of the statement is the object of the responsible stakeholder's action. The object of Example 1 is "drugstore.com member's PHI." Notice that again the rule for changing pronouns was used. The object is important because it will be the object of the commitment, right, or legal obligation.

Source. The source is the source of the object. The source will become part of the commitment, right, or legal obligation. Example 1 does not have a source because the statement does not include the source of drugstore.com member's PHI. If instead the statement had stated, "We may use your PHI received from your doctor," then the source would have been "drugstore.com member's doctor."

Target. The target of the statement is the target of the responsible stakeholder's action. The target will become part of the commitment, right, or legal obligation. There is no target in Example 1. To understand target consider the following statement, "We may disclose your PHI to third parties." For this statement the target is "third parties" because that is whom the PHI is disclosed to.

Purpose. The purpose of the statement is reason the action is performed. The purpose will explain the reason for the commitment, right, or legal obligation. The purpose for Example 1 is "treat drugstore.com member."

When listing the purposes for a statement we found that it is not necessary to classify the statement as separate commitments, rights, and legal obligations based

on multiple purposes. From the perspective of the requirements engineer, the purposes are not separate software requirements but rather constraints on the software requirements. Consider the following statement from Aetna Web Privacy Statement: "From time to time, we may request personal information from you at our sites in order to deliver requested materials to you, respond to your questions, or deliver a product or service." The right for this statement is *request personal information from you at our sites in order to (a) deliver requested materials to Aetna member, (b) respond to Aetna member's questions, or (c) deliver a product or service*, with the purposes being *(a) deliver requested materials to Aetna member, (b) respond to Aetna member's questions, or (c) deliver a product or service*. From these purposes one can see that based on these purposes additional software requirements are not needed but additional constraints, from the purposes, will exist for the software requirement. Purposes within policy documents were often preceded by "to provide."

Conditions. The conditions within the statement place restrictions on the commitment, right, or legal obligation. These describe the conditions that must be met in order for the commitment, right, or legal obligation to apply. Example 1 does not contain any conditions. If the statement had been "We may use your PHI to treat you if you have authorized us to do so," then the conditions would have been "if drugstore.com member has authorized drugstore.com" because this condition must be met before the PHI can be used to treat the drugstore.com member.

Wan and Singh state that commitments have two forms—unconditional and conditional [WS05]. Breaux and Antón also consider conditions for the obligations that they examined [BA07]. By documenting the conditions contained within a statement we can classify each statement as either a conditional or unconditional commitment, right, or legal obligation.

Scenarios. For the purpose of this paper, a scenario typically describes the ways in which a commitment, right, or legal obligation is carried out. This is in contrast to a purpose that simply describes the reason for a commitment, right, or legal obligation. In the policy documents, some commitments, rights, and legal obligations are stated then further elaborated with a concrete description. In Example 1, the sentence that begins with, "For example," clearly denoted an elaboration of the right that precedes it. Therefore, the scenario for Example 1 is

“if you are being treated for an injury, we may share your PHI with your primary physician so they can provide proper care” because this gives an example of when the PHI may be used.

When analyzing privacy documents any identified scenarios are documented in the spreadsheet. It is important to document these scenarios because they provide context that is later needed to elucidate subtle differences between stakeholders and their corresponding commitments, rights, and legal obligations.

Commitment, Right, or Legal Obligation. Depending on the type the statement was identified as, either a commitment, right, or legal obligation is listed for the statement. Each of which is formed in the following manner (assuming all of the attributes exist for that statement): [action] [object] from [source] to [target] to [purpose] given [conditions]. Since some of the attributes may be empty, the commitment, right, or legal obligation is whichever attributes exist that are in the pattern. From this pattern, the right for Example 1 is “use drugstore.com member’s PHI to treat drugstore.com member.”

4 Results and Discussion

Evaluating Aetna’s eleven policy documents revealed that the seven Notice of Privacy Practices documents were extremely similar. For the most part, these notices had many of the same commitments, rights, and legal obligations and only differed based on the plan name. As a result of this there are many duplicate commitments, rights, and legal obligations. When analyzing the policy documents we noted those commitments, rights, and legal obligations that were duplicates of statements from other documents. Table 2 summarizes the number of commitments, rights, and legal obligations extracted from the Aetna policy documents. The first column of the table denotes what the statements were classified as—commitments, rights (customer and organizational), and legal obligations. The second column has the number of unique occurrences of commitments, rights, and legal obligations. The numbers of statements that are unique conditional are shown in the third column. The fourth column has the numbers for duplicate occurrences, while the numbers for duplicate conditional occurrences are in the fifth column. The sixth column has the total counts, or the sums of the second and fourth columns; the seventh column is the totals for conditional classifications. It is important to note that the last row, total, is the

sum of the commitment, right, and legal obligation rows; the customer right and organization right rows are actually just a way to break up the right row.

The table shows that 69% of the total number of statements in the Aetna policy documents were rights—organizational rights followed closely by customer rights. Next were commitments (20%) and legal obligations (11%). In the eleven Aetna documents, we classified 142 statements as unique commitments, rights, and legal obligations, 41 of these were conditional. In total, including duplicate occurrences there were 466 statements classified and 75 were conditional. Of the total rights, the majority (56%) were organizational rights.

Table 2: Totals for Aetna Policy Documents

	Unique	Unique conditional	Duplicate occurrences	Duplicate conditional occurrences	Total	Conditional total
Commitment	56	22	37	0	93	22
Right	69	14	251	34	320	48
Customer right	31	9	110	28	141	37
Organizational right	38	5	141	6	179	11
Legal obligation	17	5	36	0	53	5
Total	142	41	324	34	466	75

Examining the commitments, rights, and legal obligations in the Aetna documents it is interesting to look at the percentages for unique compared to total; these percentages are shown in Table 3. Commitments have the highest percentage of unique occurrences within the Aetna policy documents.

Table 3: Percentage Unique in Aetna Policy Documents

	Percentage Unique
Commitment	60.2%
Right	21.6%
Customer right	22.0%
Organizational right	21.2%
Legal obligation	32.1%
Total	30.5%

This work is important because the software that software engineers create needs to comply with policy. In order to do so, software engineers need to be able to interpret the policy documents and extract requirements for the software they are creating. With policy documents, we found that an approach to extract the requirements is to first extract commitments, rights, and legal obligations from the policy documents. Looking at commitments, rights, and legal obligations lets the requirements engineer prioritize what he or she needs to ensure that the software does to comply with law. When prioritizing commitments, rights, and legal obligations, the legal obligations would have a higher importance than the commitments.

We wanted to investigate whether the rights and obligations that are expressed in law could be tied to the rights and obligations expressed in policy documents. What we learned was that the majority of the statements made in privacy policy documents are actually commitments from the organization to the customer and rights—customer and organizational, while there are fewer legal obligations present in policy documents. We learned that rights and obligations are a more appropriate unit for analysis with law than with policy documents. As we showed, privacy policy documents predominantly express commitments to customers and rights, rather than legal obligations.

5 Future Work

Our plans for future work include further investigating scenarios and purposes that are inferred from the statements within other kinds of policy documents. We also plan to explore the use of delegations within the documents to express agreements between two organizations. During our preliminary study we saw a few delegations but we need to identify a larger number to have enough to work with so that we can decide the best approach for dealing with them. We also plan to evaluate the commitments, rights, and legal obligations in additional health care privacy policy documents [AEV07].

Based on a comprehensive taxonomy of privacy commitments in health care privacy policies, we plan to design a survey instrument to investigate users' perceptions as they relate to commitments, rights, and legal obligations. Within the survey, we also plan to compare commitments, rights, and legal obligations to goals.

Acknowledgements

This work was partially funded by the following grants: NSF ITR Grant #0325269, NSF Cyber Trust Grant #0430166 and NSF Science of Design Grant #0725144. The authors would also like to thank Travis Breaux, Gurleen Kaur, Aaron Massey, Jeremy Maxwell, and Paul Otto for their comments.

References

[Ant96] Annie I. Antón, "Goal-Based Requirements Analysis," in *Proceedings of the Second IEEE International Conference on Requirements Engineering (ICRE)*, 1996, pp. 136-144.

[AEC03] Annie I. Antón, Julia B. Earp, and Ryan A. Carter, "Precluding Incongruous Behavior by Aligning Software Requirements with Security and Privacy Policies," *Information and Software Technology*, Elsevier, 45(14), pp. 967-977, 1 November 2003.

[AEV07] Annie I. Antón, Julia B. Earp, Matthew W. Vail, Neha Jain, Carrie M. Gheen, and Hack M. Frink, "HIPAA's Effect on Web Site Privacy Policies," *IEEE Security & Privacy*, January/February 2007.

[BR94] S. Gayle Baugh and Ralph M. Roberts, "Professional and organizational commitment among engineers: conflicting or complementing?" *IEEE Transactions on Engineering Management*, vol. 41, no. 2, pp.108-113, May 1994.

[Bea03] Howard Beales, The FTC's Use of Unfairness Authority: Its Rise, Fall, and Resurrection, given at the Marketing and Public Policy Conference, Washington, DC, May 30, 2003.

[BA05] Travis D. Breaux and Annie I. Antón, "Analyzing Goal Semantics for Rights, Permissions, and Obligations," in *Proceedings of the 13th IEEE International Conference on Requirements Engineering*, 2005, pp. 177-186.

[BA07] Travis D. Breaux and Annie I. Antón, "A Systematic Method for Acquiring Regulatory Requirements: A Frame-Based Approach," in *Proceedings of the 6th International Workshop on Requirements for High Assurance Systems (RHAS-6)*, 2007.

[BA08] Travis D. Breaux and Annie I. Antón, "Analyzing Regulatory Rules for Privacy and Security Requirements," IEEE Transactions on Software Engineering, Special Issue on Software Engineering for Secure Systems, 34(1), pp.5-20, January 2008.

[BAD09] Travis .D. Breaux, Annie I. Antón, and Jon Doyle "Semantic Parameterization: A Process for Modeling Domain Descriptions," ACM Transactions on Software Engineering and Methodology, April 2009 (in press).

[BVA06] Travis D. Breaux, Matthew W. Vail, and Annie I. Antón, "Towards regulatory compliance: extracting rights and obligations to align requirements with regulations," in *Proceedings of the 14th IEEE International Requirements Engineering Conference*, 2006, pp.46-55.

[Cam98] Nancy Campbell, *Writing Effective Policies and Procedures: A Step-by-step Resource for Clear Communication*. New York: AMACOM, 1998.

[CL90a] Philip R. Cohen and Hector J. Levesque, “Intention Is Choice with Commitment,” *Artificial Intelligence*, vol. 42, pp. 213-261, 1990.

[CL90b] Philip R. Cohen and Hector J. Levesque, “Persistence, Intention, and Commitment,” in *Intentions in Communication*, Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, Eds., Cambridge, MA: MIT Press, June 1990.

[CMP90] Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, Eds., *Intentions in Communication*. Cambridge, MA: MIT Press, June 1990.

[FTC14] Federal Trade Commission Act (15 U.S.C. §§ 41-58).

[GJ99] Ellen Garbarino and Mark S. Johnson, “The Different Roles of Satisfaction, Trust, and Commitment in Customer Relationships,” *The Journal of Marketing*, vol. 63, no. 2, pp.70-87, April 1999.

[GAP07a] Sepideh Ghanavati, Daniel Amyot, and Liam Peyton, “A requirements management framework for privacy compliance,” in *Proceedings of the 10th Workshop on Requirements Engineering*, 2007, 149-159.

[GAP07b] Sepideh Ghanavati, Daniel Amyot, and Liam Peyton, “Towards a Framework for Tracking Legal Compliance in Healthcare,” in *Advanced Information Systems Engineering*, J. Krogstie, A. L. Opdahl, and G. Sindre, Eds. Berlin Heidelberg: Springer-Verlag, 2007, pp. 218-232.

[GS67] B.C. Glaser and A.L. Strauss. *The Discovery of Grounded Theory*, Aldine Transaction, Chicago, IL, 1967.

[Had95] Afsaneh Haddadi, Ed., “A formal theory of commitments,” in *Communication and Cooperation in Agent Systems*, Vol. 1056, Lecture Notes in Computer Science, Berlin: Springer, 1995, 51-82.

[Hof06] Marcia Hofmann, “Federal Trade Commission Enforcement of Privacy,” in *Proskauer on privacy: a guide to privacy and data security law in the information age*, Christopher Wolf, New York, NY: Practising Law Institute, 2006.

[Jac97] Michael Jackson, “The Meaning of Requirements,” *Annals of Software Engineering*, vol. 3, Baltzer Science Publishers, 1997, pp. 5-21.

[LCN90] Hector J. Levesque, Philip R. Cohen, José H. T. Nunes, “On Acting Together,” in *Proceedings of the 8th National Conference on Artificial Intelligence*, 1990, pp. 94-99.

[Mac80] Ian Macneil, *The New Social Contract, An Inquiry into Modern Contractual Relations*. New Haven, CT: Yale University Press, 1980.

[Mag99] Dr. Timothy J. Maggs, “The Commitment Theory,” 1999. http://angelacevedo.com/drtimothymaggs.com/Articles/the%20commitment%20theory.html.

[MZD92] Christine Moorman, Gerald Zaltman, and Rohit Deshpandé (1992), “Relationships Between Providers and Users of Market Research: The Dynamics of Trust Within and Between Organizations,” *Journal of Marketing Research*, vol. 29 no. 2, pp. 314-29, August 1992.

[NYQ07] Zhang Ning-jun, Jiang Yong-zhong, and Li Qin, “Connecting Service Employee’s Organizational Commitment with Customer Perception: A Conceptual Model,” in *Proceedings of the International Conference on Service Systems and Service Management*, 2007, pp. 1-11.

[ONe02] Onora O'Neill, *Autonomy and Trust in Bioethics*. Cambridge, UK: Cambridge University Press, 2002.

[OA07] Paul N. Otto and Annie I. Antón, "Addressing Legal Requirements in Requirements Engineering," in *Proceedings of the 15th IEEE International Requirements Engineering Conference*, 2007, pp. 5-14.

[RRA06] Margaret F. Reid, Cynthia K. Riemenschneider, Myria W. Allen, and Deborah J. Armstrong, "Affective commitment in the public sector: the case of IT employees," in *Proceedings of the 2006 ACM Special Interest Group on Management Information Systems on Computer Personnel Research*, 2006, pp. 321-332.

[Rob05] William N. Robinson, "Implementing Rule-based Monitors within a Framework for Continuous Requirements Monitoring," in *Proceedings of the 38th Hawaii International Conference on System Sciences*, 2005, pp. 188a.

[Sin92] Munindar P. Singh, A Critical Examination of the Cohen-Levesque Theory of Intentions," in *Proceedings of the 10th European Conference on Artificial Intelligence*, 1992, pp.364-368.

[SA91] Munindar P. Singh and Nicholas M. Asher, "Towards a Formal Theory of Intentions," in *Proceedings of the European workshop on Logics in AI*, 1991, pp. 472-486.

[TDD99] Ayça Tarhan, Elif Demirörs, and Onur Demirörs, "A distributed tool for commitment specification and management," in Proceedings of EUROMICRO Conference, 1999, pp. 210-217 vol.2.

[WS05] Feng Wan and Munindar P. Singh, "Formalizing and Achieving Multiparty Agreements via Commitments," in *Proceedings of Autonomous Agents and Multi-Agent Systems*, 2005, pp.770-777.

[WB90] Samuel D. Warren and Louis D. Brandeis, "The right of privacy", *Harvard Law Review*, volume 4(5), pp. 193-220, 15 December 1890.

[Wu93] Paul Horng Jyh Wu, "Closing The Gap Between Discourse Structure And Communicative Intention," *Workshop On Intentionality And Structure In Discourse Relations*, 1993.