Submitted to: IEEE 2002 Symposium on Security and Privacy
NCSU CSC TR-2001-10

Goal-Mining to Examine Health Care Privacy Policies

Annie I. Antón[1], Julia B. Earp[2], Angela Reese[3]
[1] *College of Engineering, North Carolina State University, Raleigh, NC 27695-7534*
[2] *College of Management, North Carolina State University, Raleigh, NC 27695*
[3] *College of Arts and Sciences, University of Dayton, Dayton, OH 45469-0800*

[1] aianton@eos.ncsu.edu [2] Julia_Earp@ncsu.edu [3] angelamreese@hotmail.com

Abstract

Privacy has recently become a prominent issue in the context of electronic commerce Web sites. Increasingly, privacy policies posted on such Web sites are receiving considerable attention from the government and consumers. We have used goal-mining, the extraction of pre-requirements goals from post-requirements text artifacts, as a technique for analyzing privacy policies. The identified goals are useful for analyzing implicit internal conflicts within privacy policies and conflicts with the corresponding web sites and their manner of operation. These goals can be used to reconstruct the implicit requirements met by the privacy policies. We present the results of our analysis of 23 Internet privacy policies for companies in three health care industries: pharmaceutical, health insurance and online drugstores.

1 Introduction

Requirements engineering is the principled application of proven methods and tools to describe the behavior and constraints of a proposed system. As such, it arguably influences the outcome of a software project more than any other sub-discipline within software engineering [FB91] as well as the outcome of other analysis activities such as policy formation. Those of us who can offer a systems engineering perspective must assume more responsibility for approaching systems and their respective policy holistically [AEP01]. Our approach to policy and requirements specification [AE01] relies on the application of goal and scenario-driven requirements engineering methods for secure electronic commerce systems, resulting in the specification of: privacy policies, security policies and the corresponding system requirements for these proposed or envisioned systems. In this paper we explain our application of these requirements engineering techniques to Internet policy analysis. We also introduce a taxonomy of privacy goals which provides an effective mechanism for analyzing and comparing privacy policies, system requirements and the functionality of the respective systems.

Health care privacy, as it pertains to organizational practices, holds profound implications as service delivery impacts human life, legality and social policy [Dar97]. The transmission and dissemination of health care information in electronic form has raised numerous privacy concerns among both consumers and providers [EP00]. The evolving trend toward Internet supported health care services has resulted in increased information sharing among providers, pharmacies and insurers. Unfortunately, such information sharing often conflicts with consumers' desires to be shielded from unauthorized use of their personal information. In order to identify the system requirements reflected in health-care privacy policies, we employed a technique, goal-mining, to derive the privacy-related goals of various Internet health care Web sites. Our motivation is two-fold. First, we seek to develop a corpus of reusable privacy and security goals for electronic commerce software developers [AAB99, AE01]. Second, goals provide an excellent unit by which to objectively analyze and compare Internet privacy policies, enabling us to provide useful guidance to practitioners, policy makers and consumers.

The protection of personal information, such as that managed by health care Web sites, is not an option but a necessity. Goal-mining and analysis are effective techniques for examining how Internet Web sites claim they manage online customer data and how they convey these practices to their customers. Section 2 provides an overview of the state of the research and practice in health care privacy policy, policy evaluation and goal-based requirements analysis. In Section 3 we introduce our taxonomy of privacy goals and provide examples. In Section 4 we codify the specific heuristics that guide the goal-mining process, providing examples from our analysis of 23 health care Internet Web site privacy policies. Results of our analysis are discussed in Section 5. Finally, a summary and discussion of future work is provided in Section 6.

2 Background and Related Work

This section provides an overview of the relevant work in health care privacy policy, policy evaluation and the role of policy analysis in requirements engineering.

2.1 Health Care Privacy Policy

A privacy policy is a comprehensive description of a Web site's information practices; it is located in an easily accessible place on the site [FTC98, FTC00]. Every organization involved in online transactions has a responsibility to adopt and implement a policy for protecting the privacy of personally identifiable information (PII). Organizations must also foster the adoption and implementation of effective online privacy policies by their partners and subsidiaries. Although organizations engaged in electronic transactions should disclose privacy policies that are based on fair information practices [FIP73, FTC98, FTC00], the Georgetown Internet Privacy Policy Survey [Cul99] found that Internet privacy disclosures do not always reflect fair information practices. As we discuss in Section 3, this contributes to our inability to classify all privacy goals as simply protection goals. Instead, we classify privacy goals as either protection or vulnerability goals.

In 2000 over 17,000 different health care Web sites offered a wide range of products and services on the Internet [GHS00]. Examples of activities currently available on many health care Web sites include: purchasing, provision of clinical information, professional interaction and personal health records. The privacy practices of health-care related Web sites have recently received attention in the press. Eli Lilly unintentionally released email addresses of people on Prozac [Wil01]. The company sent customers an email informing them that a service was being cancelled. This email contained all the email addresses of the people currently using this service. The sensitive nature of these services and recent incidents such as this suggests a need for legislation regarding information exchange.

Although the Privacy Act of 1974 provides some protection for medical records that are held by federal agencies, it does not cover medical records held by private groups where most medical records are actually created and stored[1]. Moreover, numerous exceptions are contained in the act so that its overall protection is leaky at best. When introduced, the Privacy Act was heralded as a huge step forward, but it is now considered the "most outdated" privacy act in the world [Bar98].

The increasing utilization of the Internet for healthcare information exchange has initiated legal reform with regard to privacy protection of electronic medical records. The 1996 Health Information and Portability Accountability Act (HIPAA)[2] mandated that the U.S. Government

[1] 5 U.S.C. 552a (1994)

[2] Health Insurance Portability and Accountability Act of 1996, 42 U.S.C.A. 1320d to d-8 (West Supp. 1998).

Administration introduce regulations regarding the control of medical records. These regulations called for the inclusion of a provision for health information privacy. The Department of Health and Human Services (HHS) published the final Privacy Rule[3] that took effect on April 14, 2001, requiring health care providers and health plans to comply by April 14, 2003.

2.2 Privacy Policy Evaluation Mechanisms

Privacy policies are evaluated in a rather ad hoc and inconsistent manner. Current solutions include P3P [RC97] and various privacy seal programs [Ben99] as we now discuss.

A 1999 survey revealed that 87% of Internet users are concerned about threats to their privacy when online [CRA99]. However, several studies have subsequently shown that Internet users are more inclined to trust a Web site if it posts a privacy policy [EB02, GHS00]. These findings have produced positive results as most online companies now post some form of privacy policy on their Web site. The downside to this is that not all consumers can (or are willing to) take the time to read and understand these policies. Consequently, several privacy policy evaluation mechanisms have been introduced to assist online consumers. The World Wide Web Consortium is establishing the Platform for Privacy Preferences Project (P3P)[4] as an industry standard to provide an automated way for users to gain control over the use of their personal information on Web sites they visit. P3P requires consumers to answer a set of standardized multiple-choice questions that address various aspects of Web site privacy policies. The sites implementing P3P possess a privacy policy in machine readable format and users of these sites may configure their browsers to automatically determine if a Web site's privacy policy reflects their personal needs for privacy. This is done by comparing the user's responses to the multiple choice questions with the statements in a P3P compliant policy. As of September 2001, only thirteen sites were publicly listed as being compatible with the latest specification of P3P[5]. In addition to the slow adoption of P3P, the European Union has rejected P3P as a viable technical means for supporting their stringent privacy laws [Epi00]. A report by the Electronic Privacy Information Center (EPIC) [Epi00], asserts that P3P fails to comply with baseline standards for privacy protection and that it is a complex and confusing protocol that will hinder Internet users in protecting their privacy. There is little evidence to support the industry claim that P3P will improve user privacy, as it does nothing to assess compliance with the recommended Fair Information Practices.

Self-regulation is encouraged by the FTC [FTC98] but some online businesses still have not adopted the fundamental fair information practices that address consumer privacy. Privacy seal organizations, such as TRUSTe[6], BBBonline[7] and WebTrust[8], in a sense complicate privacy policy since consumers, understandably, often trust indirect and abbreviated indicators of privacy protection rather than reading the full privacy policy.

The seal and guarantee provided by TRUSTe appear to be quite comforting to consumers. However, most consumers are unfamiliar with what the TRUSTe privacy seal truly signifies. In reality, this privacy seal simply ensures that TRUSTe has reviewed the licensee's privacy policy for disclosure of the following uses of information by a Web site: what personal information is being gathered; how the information will be used; who the information will be shared with; the choices available regarding how collected information is used; safeguards in place to protect

---

[3] Federal Register 59918 et seq., Department of Health and Human Services, Office of the Secretary, 45 CFR Parts 160 through 164, Standards for Privacy of Individually Identifiable Health Information, (December 28, 2000).

[4] http://www.w3.org/P3P/

[5] http://www.w3.org/P3P/compliant_sites

[6] http://www.truste.com/

[7] http://www.bbbonline.com/

[8] http://www.cpawebtrust.org/

personal information from loss, misuse, or alteration; and how individuals can update or correct inaccuracies in information collected about them. This is not particularly stringent and does not reflect a real commitment to consumer privacy, merely an openness about what degree of privacy is supported. TRUSTe requires licensees to disclose their privacy practices and adhere to established privacy principles based on the fair information practices. This is an admirable service and evidence exists that it has brought about the protection of consumer privacy in a very real way in the case of Toysmart.com. However, consumers should be alarmed by the privacy policies of some web sites displaying this supposed "commitment to customer privacy." As long as a privacy policy openly admits that customer information is sold, leased, etc., a website is eligible for a TRUSTe privacy seal. For example, some TRUSTe licensees track what web page visitors were at prior to accessing their website, whereas other TRUSTe licensees sell or share their customer email lists with other companies, allowing these third parties to send you email solicitations.

The BBB*OnLine* privacy seal is posted on Web sites for which the merchant has met BBB*OnLine* Privacy Program requirements regarding notice, choice, access and security of PII collected online. These companies must post privacy policies stating what personal information is collected, how it will be used, choices consumers have in terms of use, and the policy must verify security measures taken to protect this information. These companies commit to abide by their posted privacy policies, and agree to a comprehensive independent verification by BBB*OnLine*. Similar to TRUSTe, consumers are given a false sense of security when they encounter a BBB*OnLine* seal since they do not realize that a Web site can display it regardless of whether or not a privacy policy truly protects consumer privacy.

The CPA WebTrust seal entails a more stringent privacy evaluation mechanism, as only licensed public accountants who complete special training are able to issue a WebTrust seal. Unlike other seal programs, WebTrust requires accountants to conduct an independent examination that carries the professional equivalency of a financial statement audit. A licensed CPA, Chartered Accountant, or equivalent will only award a seal to a site if it completely passes the examination. WebTrust meets the US industry consensus standards for privacy established by the Online Privacy Alliance[9]. Unlike P3P, the program substantively meets the standards for the European Union and Canada as well. However, there are very few Web sites that currently display the CPA WebTrust seal. We attribute that to the extremely high cost of a CPA WebTrust seal, especially since it is not mandatory and has not been proven to markedly boost site visits.

A more effective privacy evaluation mechanism would be a policy-rating tool that considers not only the presence of certain policy content, but the implications of the policy content in reference to how practices affect consumer privacy. Enonymous.com hosted a now defunct rating system, *www.privacyratings.com* that reviewed and rated Web site privacy policies according to how the site used PII. The three specific criteria used for the ratings were (1) whether a site contacts visitors for purposes beyond the primary purpose of data collection; (2) whether a site shares, trades, or sells user data; and (3) whether sites conduct such use with explicit visitor permission. The enonymous rating system focused on the notice and choice offered to visitors about the use of their PII. These correspond to the first two principles of the FTC's five fair information practice principles but do not include security, access/participation, enforcement/redress or other factors such as cookie use. Enonymous.com was dissolved in 2000 due to financial reasons. Given the growing concern for online personal privacy, it is evident that the public is in need of a more effective privacy evaluation mechanism. Requirements

---

[9] The Online Privacy Alliance is a group of corporations and associations having a common goal of protecting individual privacy online (http://www.privacyalliance.org/)

engineering provides reliable and straightforward mechanisms for evaluating privacy as we now discuss.

2.3 Policy from the Requirements Engineering Perspective

Although researchers in the requirements engineering community are beginning to focus on electronic commerce applications [AP98, ACD01, Rob97] there remains a need to apply proven requirements analysis methods (a routine activity in software engineering) and demonstrate how to best apply these methods within the context of establishing and analyzing policy. Goal analysis and scenario analysis have been successfully applied within the context of evolving electronic commerce systems [AP98] as we now discuss.

2.3.1 Goals

Goals are the objectives and targets of achievement for a system. In requirements engineering, goal-driven approaches focus on why systems are constructed, expressing the rationale and justification for the proposed system [Lam01]. Since goals are evolutionary, they provide a common language for analysts and stakeholders. Focusing on goals, instead of specific requirements, allows analysts to communicate with stakeholders using a language based on concepts with which they are both comfortable and familiar. Furthermore, since goals are typically more stable than requirements [Ant97], they are a beneficial source for requirements derivation. Goals are operationalized and refined into requirements and point to new, previously unconsidered scenarios.

2.3.2 Goal-Based Requirements Engineering

The Goal-Based Requirements Analysis Method (GBRAM) [Ant96, Ant97, AP98, ADS00] is a straightforward methodical approach to identify system and enterprise goals and requirements. It is useful for identifying and refining the goals that software systems must achieve, managing trade-offs among the goals, and converting them into operational requirements. The method suggests goal identification and refinement strategies and techniques through the inclusion of a set of heuristics, guidelines and recurring question types. Five sets of heuristics are included: identification heuristics, classification heuristics, refinement heuristics, elaboration heuristics and conflict identification/resolution heuristics. The heuristics are useful for identifying and analyzing specified goals and scenarios as well as for refining these goals and scenarios. The GBRAM heuristics and supporting inquiry include references to appropriate construction of scenarios and the process by which they should be discussed and analyzed. The method has been successfully applied to the analysis of electronic commerce applications [AP98, ACD01]. In this paper, we describe our use of the method to mine privacy policies for system goals and requirements and codify the domain specific heuristics for applying the GBRAM for goal-mining Internet privacy policies. In the following sections we introduce our privacy goal taxonomy and describe the goal-mining process.

3 Taxonomy of Privacy Goals

During the summer of 2000, we engaged in a *goal-mining* exercise in which we evaluated 24 Internet Privacy Policies from 8 non-regulated electronic commerce industries (e.g. Online travel agencies and online retailers). The identified goals are useful for discovering implicit internal conflicts within privacy policies and conflicts with the corresponding web sites and their manner of operation. These goals can also be used to reconstruct the implicit requirements met by the privacy policies and can also be used to reason about expected privacy policy content for different types of Web sites. This information can assist developers in creating policies that address common goals for a given Web site type.

Initially, privacy experts who had viewed our pilot study data suggested that all goals expressed in a Web site's privacy policy should support the Code for Fair Information Practices [FIP73]. However, the goals derived from 24 Internet electronic commerce privacy policies proved challenging to classify in this simple manner. We attempted to classify the derived privacy goals according to the Fair Information Practices (notice / awareness; choice / consent; access / participation; integrity / security; and enforcement / redress). Unfortunately, it became abundantly clear that it was impossible to "force-fit" all the derived goals into the five FIP goal types. We thus analyzed the remaining unclassified goals to determine what was different about those goals than the goals that supported the fair information practices. Careful examination revealed that the remaining goals did not exemplify privacy protection practices; instead, they reflected practices that introduce vulnerabilities in a site's ability to protect personal information. This led us to create a taxonomy for privacy-related system goals so that consumers and system developers can more accurately compare privacy practices and reason about a site's functionality and alignment with privacy policies.

In our taxonomy, we broadly classify privacy goals as either privacy protection or privacy vulnerability goals. *Privacy protection goals* are those that relate to the five FIP principles and are related to the desired protection of consumer privacy rights. *Privacy vulnerability goals* are those related to existing threats to consumer privacy. In contrast to protection goals, privacy vulnerability goals are those that represent statements of fact or existing behavior and are often characterized by privacy invasions. We define the five kinds of *privacy protection goals* in Table 1 and the five kinds of *privacy vulnerability goals* in Table 2. These goal classes are useful for analyzing implicit internal conflicts within privacy policies and conflicts with the corresponding Web sites and their manner of operation. Goals can be used to reconstruct the implicit requirements met by the privacy policies.

Table 1: Privacy Protection Goal Taxonomy Classifications

Protection Goal Taxonomy	Protection Goal Sub-Classifications
Notice/Awareness Goals asserting that consumers should be notified and/or made aware of an organization's information practices before any information is actually collected from them (e.g., an organization's privacy policy).	• General Notice/Awareness • Identification of the Uses to Which the Data Will be Put • Identification of Any Potential Recipients of the Data • 3rd Party Limitations • Nature of the Data Collected • Steps Taken by the Data Collector to Ensure the Confidentiality, Integrity, & Quality of the Data
Choice/Consent Goals ensuring that consumers are given the option to decide what personal information collected about them is to be used and whether it may be used for secondary purposes.	• Choice of How Data is Used • Choice of Sharing Data • Choice of What Data is Taken/Stored
Access/Participation Goals allowing or restricting access to a particular site or functionality based on whether or not the consumer provides their PII.	• PII Provision Required • PII Provision Optional
Integrity/Security Goals ensuring that data is both accurate and secure. Security and accuracy comes from both the consumer and the organization collecting the PII. Goals in this category range from vague statements stating only that PII is kept securely to specific technical goals of what security protocols will be used to transfer PII over the Internet.	• Mission Statement • User-Supplied Integrity Goals • Using Anonymous PII • Providing Consumer Access to Data • Destroying Untimely or Sensitive Data • Managerial Measures to Protect Against Loss and the Unauthorized Access, Destruction, Use, or Disclosure of the Data • Technical Measures to Protect Against Loss and the Unauthorized Access, Destruction, Use, or Disclosure of the Data
Enforcement/Redress Goals addressing the mechanisms that are in place to enforce privacy, otherwise a policy is merely suggestive, rather than prescriptive. Prescribe a way of working and general guidelines companies should follow. These include both self-imposed and government imposed work restrictions.	• Operational Prevention Assurance • 3rd Party Prevention Assurance • Failure of Assurance

Table 2: Privacy Vulnerability Goal Taxonomy Classifications

Vulnerability Goal Taxonomy	Vulnerability Goal Sub- Classifications
Information Monitoring Goals concerning what organizations may track what consumers do on their site through means such as cookies. This could be for the consumer's benefit, like when an electronic-commerce application maintains a shopping cart for a consumer, or for the organization's benefit, be it for purely statistical use or for profit (via selling of aggregated information to 3rd parties).	• Monitoring for Personalization • Monitoring for Statistics • Limitation of Monitoring
Information Aggregation Aggregation combines previously gathered PII data with data from other sources, furthering their own PII regardless of whether the user has provided such data or not.	N/A
Information Storage Goals addressing how and what records are stored in an organization's database. These goals cover a broad range, from security to monitoring and basically storage-specific.	• Storage for Customer Use • Storage for Corporate Use
Information Transfer Goals concerning any transfer of information. Privacy by its very definition means an insurance that others can not find something out. This wholly incorporates the idea that information must not be transferred. These goals address safeguards against the transfer of information, as well as to whom what information is transferred.	• Sharing PII with users • Sharing/Selling with Other Companies/Sites • Limitation of Sharing
Information Collection Goals addressing the how and what information is being collected. Collection occurs when an organization collects information from a consumer either by way of directly asking them to put the information in, or by collecting information without their consent, such as browser information.	• Direct Collection (e.g. user provided information) • Indirect Collection (e.g. browsing patterns)
Information Personalization Goals addressing personalization as when consumers either change their PII, thus affecting the functionality or content offered to them.	• Personalization by User Preference • Personalization of Site and Service • Personalization of Advertising, Offers, and Promotions
Contact These goals deal with how and for what purpose organizations contact consumers using their PII. This could be helpful, such as contacting customers to validate an email address, or annoying, such as sending out unwanted promotions based on past patterns.	• Contact for Promotions and Offers • Contact for Security and Verification • Contact Based on Preference

Table 3 provides concrete examples of privacy protection goals and privacy vulnerability goals. Eventually, in software development, these goals are operationalized into system requirements and checked for compliance with the respective policies. Our preliminary analysis showed that several Web sites with these goals stated in the privacy policies do not actually comply with the goals, a subject of discussion for a future paper.

Table 3: Sample Protection and Vulnerability Privacy Goals

Privacy Protection Goals	Privacy Vulnerability Goals
Notice/Awareness NOTIFY users before data is collected NOTIFY users of updates to privacy policy **Choice/Consent** ALLOW customer to opt-in to sharing PII w/ member sites OPT-IN to controlling whether to have PII stored **Access/Participation** ALLOW customer to check their PII for accuracy ALLOW customer to modify their PII **Integrity/Security** CROSS-REFERENCE user info to find uses of multiple IDs or aliases **Enforcement/Redress** DISCIPLINE associates/employees who violate privacy policy	**Information Collection** COLLECT children's names and ages when they enter contests COLLECT user browsing patterns **Information Monitoring** MONITOR customer site usage patterns **Personalization** CUSTOMIZE offers based on customer's account and purchase records **Information Storage** STORE purchase records **Information Aggregation** AGGREGATE purchase info by zip code **Information Transfer** SHARE PII w/ third parties

4 The Goal Mining Process

In this section, we present the goal-mining process and its associated heuristics within the context of our analysis of Internet health care privacy policies. The process of identifying high-level goals is fundamental to the requirements analysis and specification process. *Goal mining* refers to the extraction of goals from data sources (in this case, privacy policies) by the application of goal-based requirements analysis methods [Ant97]. The extracted goals are expressed in structured natural language, as shown in Table 3. Analysts begin the goal-mining process by first exploring any available information sources such as existing security and privacy policies, or requirements specifications and design documentation, to identify both strategic and tactical goals. Strategic goals are those that reflect high-level enterprise goals whereas tactical goals involve short-term goal achievement [AE01, PFI99]. These goals are documented and annotated with auxiliary information including the responsible agents. Goals are then organized according to goal class (privacy protection or privacy vulnerability, as previously discussed) as well as according to keyword (examples of keywords are shown in uppercase in Table 3) and subject (e.g. browsing patterns, personalization, cookies, etc. in Table 5).

Detailed techniques and heuristics for each of these operations are described in two theses [Ant97, Dem00]. Once goals are identified, they are elaborated; goal elaboration entails analyzing each goal for the purpose of documenting goal obstacles, scenarios, constraints, pre-conditions, post-conditions, questions and rationale. Goal refinement consists of removing synonymous and redundant goals, resolving any inconsistencies that exist within the goal set, and operationalizing the goals into a requirements specification. To date, we have not prepared a requirements specification. Instead, our objective has been to create a library of reusable security and privacy goals; and we are well on our way to achieving this objective. The goal library will enable developers to operationalize those goals required by their respective systems. The availability of this library of privacy and security goals in the *SMaRT* (Scenario Management and Requirements Tool) [AAB99] will enable requirements engineers and analysts to begin to build security and privacy into the e-Commerce applications early on rather than having to add it in afterwards due to oversight or external pressures.

We examined three kinds of health care Web site privacy policies using our goal-mining heuristics and privacy goal taxonomy to aid our analysis. Goals were extracted from a total of 23 privacy policies: 6 pharmaceutical companies, 7 health insurance companies, and 10 online pharmacies. The Web sites of the specific policies analyzed for each of these three health care sectors are listed in Table 4. The goals were analyzed according to different characteristics such as protection vs. vulnerability goals and subject matter (e.g. cookies, PII, browsing patterns, etc).

Our pilot study led to the development of our privacy goal taxonomy and enabled us to codify a comprehensive set of goal-mining heuristics tailored to the analysis of privacy policies, as discussed in this section. The goal-mining process is comprised of three main activities: goal identification, classification and refinement. The heuristics to guide the goal-mining process are codified below. These heuristics are broadly applicable and are not simply relevant for analysis of privacy and/or security policies. We now provide a brief overview of some of the most useful heuristics.

Heuristics for Identifying Goals

To identify goals, each statement in a privacy policy is analyzed by asking, *"What goal(s) does this statement or fragment exemplify?"* and/or *"What goal(s) does this statement obstruct or thwart?"* The identified goals are worded to express a state that is true, or the condition that holds true, when the goal is realized. Consider Privacy Policy #1 from the Blue Cross Blue Shield (BCBS) privacy policy:

Privacy Policy #1: *Our cookies will never be used to track your activity on any third party web sites or to send spam, …*

By asking these goal identification questions, we identify the goals: G_1: `PREVENT cookies from tracking activity on other websites` and G_2: `PREVENT use of cookies to send spam`.

Table 4: Privacy Policies Analyzed

	Company Name	Privacy Policy URL	Number of Protection Goals	Number of Vulnerability Goals
Health Insurance	AETNA	http://www.aetna.com/privacy.htm	5	5
	AFLAC	http://www.aflac.com/	1	1
	BCBS	http://www.bcbs.com/legal/privacystatement.html	13	7
	CIGNA	http://www.cigna.com/general/misc/privacy.html	6	5
	EHealthInsurance	http://www.ehealthinsurance.com/ehealthinsurance/PrivacyAndSecurity.html	6	9
	Kaiser Permanente	http://www.kaiserpermanente.org/disclaimer.html	3	2
	OnlineHealthPlan	http://www.onlinehealthplan.com/oasys/	7	10
Online Drugstore	CornerDrugstore	http://www.cornerdrugstore.com/store/home/about_security_privacy.asp	13	11
	DestinationRX	http://www.destinationrx.com/policy.asp	14	20
	Drugstore	http://www.drugstore.com/cat/11867/tmpl/default.asp?catid=15729	13	16
	Eckerd	http://www.eckerd.com/promise/legal_promises.asp	8	7
	HealthAllies	http://www.healthallies.com/privacy/privacy.jhtml	9	8
	HealthCentral	http://www.healthcentral.com/terms/privacypolicy/privacypolicy.cfm	11	14
	IVillage	http://www.ivillage.com/help/privacy.html	19	20
	PrescriptionOnline	http://www.prescriptiononline.com/?p=privacy	8	5
	PrescriptionsByMail	http://www.prescriptionsbymail.com/Info/Privacy.asp	10	8
	WebRX	http://www.webrx.com/	16	9
Pharmaceutical	Bayer	http://www.bayercare.com/privacy.html	7	10
	Glaxo Wellcome	http://www.imgw.com/privacy.htm	4	8
	Lilly (Eli)	http://www.lilly.com/privacy.html	2	5
	Novartis (Ciba)	http://www.ciba.com/popups/fr_legal_disclaimer_right_privacy.html	16	7
	Pfizer	http://www.pfizer.com/privacyfrm.html	4	3
	Pharmacia & Upjohn	http://www.pnu.com/Privacy/Privacy.asp	9	9

All action words are possible candidates for system goals. Goals in privacy policies may thus also be identified by looking for useful keywords (verbs). This is an extension of previously supported techniques [Abb83, RBP91, Boo91]. The following list of keywords are commonly found in most privacy policies: `ADVISE`, `AGGREGATE`, `ALLOW`, `COLLECT`, `COMPLY`, `CUSTOMIZE`, `DISALLOW`, `DISCIPLINE`, `DISCLOSE`, `ENSURE`, `IMPROVE`, `KEEP`, `LIMIT`, `MAINTAIN`, `MONITOR`, `NOTIFY`, `OPT-IN`, `OPT-OUT`, `PREVENT`, `PROHIBIT`, `PROTECT`, `PROVIDE`, `RECOGNIZE`, `REMOVE`, `REPORT`, `REQUIRE`, `RETRIEVE`, `SELL`, `SEND`, `SHARE`, `STORE`, `TRACK`, `TRANSMIT`, `TRANSFER`, and `USE`. To demonstrate the action word approach, consider the following privacy policy statement from the Eckerd Drugs privacy policy:

> **Privacy Policy #2:** *Examples of information collected include the kind of web browser you used, the domain from which you connected to the Internet, the date and time you accessed the site, your computer's operating system, and the address of the website from which you connected to our site.*

The action word `COLLECT` appears in Privacy Policy #2. This action word serves as an indicator for several goals: G_3: `COLLECT browser type`, G_4: `COLLECT domain name`, G_5: `COLLECT operating system`, G_6: `COLLECT date and time site was accessed`, and G_7: `COLLECT address of preceding website`. Goals are thus also identified using inquiry-driven and traditional action word location techniques.

Although not detailed in this paper, we also have heuristics that suggest synonymous words that may be expressed using one of the previously listed goal keywords. For example, consider Policy Policy #3, taken from the Bayer privacy policy.

> **Privacy Policy #3:** *We use the information from cookies to provide services better tailored to our users' needs and we never save passwords or credit card information in cookies.*

In this privacy policy, the term "tailor" is clearly synonymous with "customize". Using our heuristics, which guide the identification and mapping of synonymous words to our list of acceptable keywords, we express the goal G_8: `CUSTOMIZE experience at our site using cookies`. This goal, although expressed differently on different sites, appeared in 10 of the 23 analyzed privacy policies.

Heuristics for Classifying Goals

Classification of goals involves differentiating goals according to goal class (e.g. protection vs. vulnerability) and subject matter.

Protection goals are classified by analyzing each goal and asking, *"Does this goal protect one's private information?" and/or "Does this goal support one of the five fair information practices?"* Whereas, vulnerability goals are classified by considering each goal and asking *"Does this goal potentially compromise the privacy and/or security of one's private information?"* and/or *"Does this goal conflict with any of the five fair information practices?"* Consider Privacy Policy #1, which yielded the goal G_2: `PREVENT use of cookies to send spam`, this goal clearly seeks to protect one's privacy and is thus classified as a privacy protection goal. In contrast, the HealthCentral goal, G_9: `ALLOW 3rd parties to collect browsing and usage patterns information`, is a privacy vulnerability goal.

Table 5 provides an overview of the subject matter analysis. The 13 subject matters studied are listed in the left most column of the table. This part of our analysis is clearly domain specific; for example, PII/HI refers to Personally Identifiable Information and Health Information (as in medical records concerning one's prescription medication, etc.). However, we find it useful to reason about the subject matter of a particular policy since one would clearly not expect certain subjects to appear in every Internet privacy policy. We observed both privacy protection and vulnerability goals within each of the subject matter categories. This analysis is discussed in more detail in Section 5. The table details additional data about the identified goals, according to subject matter, such as the number of functional, operational, synonymous, redundant and final goals; we now discuss this refinement process.

Goal Refinement Heuristics

Organization of goals entails eliminating redundancies and reconciling synonymous goals. Goals are considered synonymous if their intended states are equivalent or if they mean the same thing to different stakeholders who simply express the goal using different terminology. It is up to the analyst to identify these instances. For example, the goals `<TRACK pages on our site using cookies>` and `<TRACK usage patterns using cookies>` are synonymous and can be reconciled as one goal which encompasses the spirit and scope of both. The analyst can choose either of the two goal names; however, all related essential information must be maintained. In the case of these two goals, they were further merged with another goal: `<TRACK usage patterns using aggregate data>`. The previous two goals were merged with the latter as follows: G_{10}: `TRACK usage patterns (using aggregate data or cookies)`. Thus, if the same goal appears more than once, all but one of the goals should be eliminated. In Table 5, merged goals are represented by the number that appears within parentheses, following the number of synonymous goals. Redundancies and synonymous goals are most easily identified after the goals have been organized according to subject matter.

Table 5: Subject Matter Goal Classes

Subject Matter	Total	Functional	Operational	Synonymous	Redundant	Final	% Reduction
Cookies/Bugs	14	7			1	7	50
Browsing Patterns/Site Usage	16			8 (1)		6	62.5
IP Address	4			1		3	25
Aggregate Info	12	3		1 (1)		7	41.7
Information	18			1 (1)		15	17
PII/HI	49	1		8 (2)	10	26	47
PII/HI Usage	42	1		13 (6)	8	14	67
Credit Card Info	9			1 (1)	3	4	56
Policies/Procedures	29	5	6	3		15	48
Contacting Customer	14		1	1	6	5	64
OPT In/Out	10			1		9	10
Security / Access	33	3	1	13 (1)	3	12	64
Children	13		1	2	2	8	38
TOTAL	263	20	9	53 (13)	33	131	50.2

Table 5 shows the number of goals that were deemed synonymous or redundant in our analysis of health care privacy policies. When reducing the number of goals, the *Browsing Patterns/Site Usage, PII/HI Usage, Contacting Customer and Security/Access* goal subjects enjoyed the greatest reduction rate. This indicates a tendency for Web site privacy policies to over-explain these particular practices using redundant/synonymous goals or statements.

The "Total" and "% Reduction" columns in Table 5 characterize the evolution of the goal set, showing the growth and refinement of the goal set throughout the goal-mining process. Our raw data initially contained 263 goals, mined from the 23 privacy policies; upon the completion of the goal refinement activities, the goal set had been reduced to 171 goals. Some goals were not truly relevant to privacy or privacy-related functionality. These goals were classified as either functional (meaning they support some system features) or operational (these goals represent business rules or operational procedures). The goal `<AGGREGATE survey results>` is an example of a functional goal; the goal `<REVIEW web security weekly>` is an example of an operational goal.

Our privacy goal taxonomy and goal-mining heuristics were validated throughout the course of our health care privacy policy analysis case study. The following section details our observations and findings.

5 Observations and Discussion

This study had several objectives, to: (1) create a taxonomy for classifying privacy goals; (2) develop a set of reusable privacy and security goals for electronic commerce software developers; and (3) use those goals to analyze and compare Internet privacy policies. This section discusses the statistical analysis that followed the goal analysis of the 23 health care privacy policies discussed in Section 4. Comparing these policies using goals is an innovative and effective analysis method that enables us to provide useful guidance to practitioners, policy makers and consumers.

5.1 Data Analysis

Most Web sites today display privacy policies that describe the site's privacy related information practices. The FTC recommends that these policies focus on practices that revolve around the five FIPs [FTC98, FTC00]. However, in spite of the many guidelines for the content and layout of these policies, privacy policy content inevitably differs from site to site. As one would expect, a site that supports electronic commerce transactions will obviously require more policy statements that focus on PII related privacy. The subject matter goals one expects to see in these site's policies include credit card information, PII, information transfer and storage. In contrast, sites whose primary mission is information dissemination with few transactions have little or no need to address the use of credit card information. This is one of the many reasons that privacy policies are so difficult to compare without consideration of the domain, business, and system requirements. It is also why goals and the goal taxonomy presented in Section 3 provide such an effective unit for measuring and comparing these policies.

The primary variables of interest for the analyses were the kind of Web site (pharmaceutical, drugstore, insurance), subject matter (browsing patterns, IP address, PII, etc.) and protection/vulnerability goals. Prior to our data analysis, we set forth several tentative assumptions in order to draw out and test their logical or empirical consequences. These hypotheses were tested using analysis of variance (ANOVA) and t-tests for correlation and paired observations. The rejection criterion for the overall test of significance in the ANOVA and t-tests was set at 0.01. Our hypotheses and conclusions for each are summarized in Table 6.

Table 6: Hypotheses

Hypotheses	Conclusion
Hypothesis 1. The number of goals in a health care privacy policy will depend on the type of site that posts that policy.	Supported
Hypothesis 2. The number of goals in a health care privacy policy will depend on the subject matter of that policy.	Supported
Hypothesis 3. The number of privacy goals in a health care privacy policy is positively associated with the word count of that policy.	Supported
Hypothesis 4. The proportion of protection goals in a health care privacy policy is greater than the proportion of vulnerability goals for that policy.	Not Supported

5.1.1 Type of Site

The first hypothesis stated that the number of privacy goals for a health care privacy policy would depend on the type of site that posted that policy. As expected, we found highly significant ($p<0.0001$) differences among the three types of health care Web sites (Health Insurance, Drugstore and Pharmaceutical). This means that the number of goals in a privacy policy will depend on the type of Web site, thus supporting Hypothesis 1. Exploring this relationship in more detail, we observe that online drugstores generally require more goals in their privacy policies when compared to health insurance and pharmaceutical policies. We assume this to be due to the transactional nature of their business. Similar to drugstore sites that allow online purchasing, health insurance sites oftentimes allow consumers to access and modify their personal information online (e.g. http://www.cigna.com). However, health insurance sites tend to require the least number of goals. We attribute this to the fact that health insurance companies are more regulated and thus, have less flexibility in how they manage personal information. Additionally, consumers tend to have less apprehension regarding the information practices in a non-retail industry [EB02].

5.1.2 Subject Matter Type

As established in Section 4, the thirteen subject matter goal categories effectively categorize the relevant content of a given privacy policy. However, it is not necessary for all privacy policies to address all thirteen subject matter categories; it depends upon the nature and mission of the given Web site. For example, an information dissemination Web site that does not execute transactions will probably have no reason to address credit card information in its policy.

Among the thirteen subject matter categories, certain privacy related issues require more goals than others do. Consider a Web site that collects only one piece of personal information, visitor IP address for the purpose of diagnosing problems with the server and administering the site. Such a practice is very straightforward and does not compromise most visitors' personal privacy. In contrast, consider a Web site that supports a great deal of transaction-oriented traffic, requiring the collection and storage of credit card numbers associated with other PII. The privacy policy for the latter Web site will need to include more goals to provide greater coverage of their information collection practices since credit card data is obviously more sensitive than IP address. This suggests a second factor, subject matter type that can determine the number of policy goals.

Hypothesis 2 stated that the number of goals in a health care privacy policy would depend on the subject matter of that policy. As expected, we found highly significant ($p<0.0001$) differences among the 13 subject matter goals (cookies, site usage, etc., see Table 5). This means that the number of goals in a privacy policy will depend on the subject matter goals included in that policy, thus supporting Hypothesis 2. Exploring these differences individually we observed that more goals would be present in a privacy policy that addressed *PII/HI*, *PII/HI Usage Cookies/Bugs or Policies/Procedures* when compared to the other subject matter types. Consider a simplified example of a policy that discusses only IP addresses; it will have fewer goals than a policy that discusses only PII.

The collection and storage of credit card information is one of the primary concerns of consumers who engage in online transactions [EB02]. Therefore, we expected Web sites to include a large number of goals disclosing practices regarding credit card data. Surprisingly, the *Credit Card Information* goal class did not factor into this grouping of goal categories with a heavy goal concentration. This could be due to an additional unknown variable that we have yet to consider in this analysis since we found no statistically significant interaction effect between subject matter and type of site.

We also observed that health insurance site policies rarely include goals regarding children. Only one health insurance site, BCBS, in our study had a goal related to children. This was expected, as children have little reason to visit insurance sites. However, COPPA[10] (the Children's Online Privacy Protection Act), which went into effect April 21, 2000, states that a company's Web site privacy policy must reveal the procedure to seek verifiable consent from a parent and the responsibilities an operator has to protect children's privacy and safety online. It is the only act that pertains specifically to Web sites. In this particular case, our goal analysis

[10] Federal Register 59887 et seq., Federal Trade Commission, 16 CFR Part 312, Children's Online Privacy Protection Act (November 3, 1999).

enabled us to determine that 22 health care Web sites have yet to address their children-related information practices as mandated by the COPPA.

5.1.3 Word Count

Within the 23 policies analyzed in this study, the word count ranged from 44 words in the AFLAC (Health Insurance) privacy policy to 3,656 words in the iVillage (Drugstore) policy. A logical presumption from this is that the policies having a higher word count will express a greater number of goals. Hypothesis 3 proposed that the number of goals in a health care privacy policy would be positively associated with the word count of that policy. As expected, our sample revealed a positive correlation, $r = 0.92$, and supported our hypothesis with $p<0.0001$. Although some consumers may find it difficult to parse the longer privacy policies, this finding suggests that the high word count might be necessary for expressing the relevant goals in certain privacy policies. Furthermore, $r^2 = 0.85$ tells us that 85% of the total variation in number of goals between Web sites is explained by the linear regression of number of goals on word count. It is important to note that there may be additional variables that may influence word count, such as whether the policy author is a lawyer (which is often the case) or due to the inclusion of definitions in policy preface material. However, consumers should benefit from the fact that privacy policies tend to express an appropriate number of goals, thus eliminating the notion that long privacy policies are excessively and unnecessarily wordy.

5.1.4 Protection and Vulnerability Goals

The five FIPs oblige Web sites to address notice/awareness, access/participation, choice/consent, integrity/security and enforcement/redress in their site privacy policies. However, as noted in section 3, it is impossible for Web sites to completely disclose their information-related practices (see definition of privacy policy in section 2.1) by simply addressing the FIPs. Therefore, during our goal-mining effort we observed that policies tend to contain both privacy protection and vulnerability goals. Hypothesis 4 stated that the proportion of protection goals in a health care privacy policy would be greater than the proportion of vulnerability goals for that policy. When comparing the number of protection goals to the number of vulnerability goals for each Web site, our analysis revealed no significant difference between them. In other words, we observed the number of protection goals for a given Web site to be equal to the number of vulnerability goals in that Web site. For example, as was the case with most of the Web sites in our sample, AETNA's privacy policy stated five vulnerability goals and five protection goals. This finding is noteworthy (and possibly even alarming) for consumers who hope that a health-care Web site would focus more on expressing how they protect their customers personal information, but that is not the case. Having an equal number of vulnerability goals demonstrates that Web sites continue to introduce risk to its customers.

5.2 Qualitative Observations

A few interesting observations surface when examining the interaction of the protection / vulnerability goal types with the subject matter goal classes. In our sample, none of the sites had protection goals in the Browsing Patterns / Site Usage and IP Address categories. This implies that these sites do not deem it necessary to protect visitor browsing patterns or IP Addresses. Similarly, no Health Insurance or Pharmaceutical sites had any Vulnerability goals pertaining to Opting In or Out. From this, we can infer that the opt-in and opt-out options at these kinds of sites are favorable to consumers as they are described with protection goals. All credit card goals (which appear only in drugstore and pharmaceutical sites) were protection goals, suggesting that a person's credit card information is considered more "important" and therefore protected as "highly private" information.

The fact that requirements specifications are often incomplete also applies to privacy policies. A careful analysis of selected goals revealed that one privacy policy failed to include the goal <`ALLOW third parties to use cookies`> even though the respective Web sites do in fact allow cookies to be sent to third parties. By setting our Netscape Preferences to accept all cookies and warn before accepting a cookie, we were able to test those sites that specifically fail to include any mention of cookies sent back to the third parties. Drugstore.com requires cookies to be enabled before a visitor may even view their home page; moreover, once cookies are enabled, this Web site sends cookies to third parties, and yet this was not expressed in their privacy policy.

We observed that privacy vulnerability goals signal potential privacy invasions. Some invasions are insidious or covert in that they are not readily apparent to consumers, as is often the case when non-transient cookies are placed on a consumer's hard drive. This is especially true when the cookie ads provide no additional value or benefit to the consumer, such as with cookies which offer personalization or purchase history information. Alternatively, some privacy invasions are obvious in that the consumer is aware or becomes aware of the privacy invasion, such as when a consumer begins to receive solicitations via email. Finally, some privacy invasions are benign; that is to say the consumer is a knowing and active contributor, facilitator, or participant in the exchange of personally identifiable information. It should be noted that what one consumer considers a privacy invasion may be a valued feature or service to another consumer. This debate is outside the scope of this paper; however, we are currently creating a privacy values survey instrument to assess these value differences and create a privacy values baseline.

6 Summary and Future Work

A company's Web site must reflect its privacy policy, else that policy is meaningless. A privacy policy describes the kinds of information collected by the Web site and the way that information is handled, stored, and used; if the Web site does not conform to its policy, the company may be subject to public outcry or legal action. In this paper, we introduce a taxonomy for classifying privacy goals; we describe our use of a software requirements engineering technique, goal-mining, to examine privacy policies for system goals and requirements; and we codify domain specific heuristics for applying the GBRAM for goal-mining Internet privacy policies. While we emphasize privacy policy goal-mining in this paper, the techniques we have presented are generalizable to different software systems; for example, security goals may be observed in most multi-user systems. Examining and comparing privacy policies using goals is an innovative and effective analysis method that enables us to provide useful guidance to practitioners, policy makers and consumers. Our plans for future work include development of a privacy rating tool based on the goal analysis process and the values baseline that will be established using our, previously mentioned, privacy values survey instrument.

Acknowledgements

This work was partially supported by the Computing Research Association's Distributed Mentor Project during the summers of 2000 and 2001 as well as NSF ITR Grant #0113792. The authors wish to thank Kevin Farmer, Colin Potts, Hema Srikanth and Ha To. Additionally, we thank Thomas Alspaugh, Richard Smith and Gene Spafford for discussions leading to our classification of privacy protection and vulnerability goals.

References

[AAB99] T. Alspaugh, A.I. Antón, T. Barnes and B. Mott. An Integrated Scenario Management

Strategy, *IEEE Fourth International Symposium on Requirements Engineering (RE'99),* University of Limerick, Ireland, pp. 142-149, 7-11 June 1999.

[Abb83] R.J. Abbot. Program Design by Informal English Descriptions. *Communications of the ACM,* 26(11):882-894, November 1983.

[ACD01] A.I. Antón, R.A. Carter, A. Dagnino, J.H. Dempster and D.F. Siege. Deriving Goals from a Use-Case Based Requirements Specification, *Requirements Engineering Journal*, Springer-Verlag, Volume 6, pp. 63-73, May 2001.

[AE01] A.I. Antón and J.B. Earp. Strategies for Developing Policies and Requirements for Secure Electronic Commerce Systems. in *E-Commerce Security and Privacy*, edited by Anup K. Ghosh, Kluwer Academic Publishers, pp. 29-46, 2001.

[AEP01] A.I. Antón, J.B. Earp, C. Potts and T.A. Alspaugh. The Role of Policy and Privacy Values in Requirements Engineering, *IEEE 5th International Symposium on Requirements Engineering (RE'01),* Toronto, Canada, pp. 138-145, 27-31 August 2001.

[AMP94] A.I. Antón, W.M. McCracken and C. Potts. Goal Decomposition and Scenario Analysis in Business Process Reengineering, *Advanced Information System Engineering: 6th International Conference, CAiSE '94 Proceedings*, Utrecht, The Netherlands, 6-10 June 1994, pp. 94-104, 1994.

[Ant96] A.I. Antón. Goal-Based Requirements Analysis, *Second IEEE International Conference on Requirements Engineering (ICRE '96),* Colorado Springs, Colorado, pp. 136-144, 15-18 April 1996. [Ant97] A.I. Antón. *Goal Identification and Refinement in the Specification of Software-Based Information Systems*, Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, GA, 1997.

[Ant97] A. I. Antón. *Goal Identification and Refinement in the Specification of Software-Based Information Systems*, Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, GA, 1997.

[AP98] A.I. Antón and C. Potts. The Use of Goals to Surface Requirements for Evolving Systems, *International Conference on Software Engineering (ICSE '98),* Kyoto, Japan, pp. 157-166, 19-25 April 1998.

[Bar98] Bartley L. Barefoot, Enacting a Heath Information Confidentiality Law: Can Congress Beat the Deadline?, 77 *N.C.L. Rev*. 283 (1998).

[Ben99] P. Benessi TRUSTe: An Online Privacy Seal Program. *Communications of the ACM*. 42(2), pp.56 – 59. February 1999.

[BEP00] D. Baumer, J.B. Earp and F.C. Payton. Privacy of Medical Records: IT Implications of HIPAA. *ACM Computers and Society,* 30(4), pp.40-47, December 2000.

[Boo91] G. Booch. *Object-Oriented Design with Applications*. Benjamin Cummings. Redwood City, California, 1991.

[CRA99] L.F. Cranor, J. Reagle and M.S. Ackerman. Beyond Concern: Understanding Net Users' Attitudes About Online Privacy, *AT&T Labs-Research Technical Report TR 99.4.3,* http://www.research.att.com/library/trs/TRs/ 99/99.4/99.43/report.htm, April 1999.

[Cul99] M.J.Culnan, Georgetown Internet Privacy Policy Survey: Report to the Federal Trade Commission. Washington, DC: Georgetown University, The McDonough School of Business, http://www.msb.edu/faculty/culnanm/gippshome.html, 1999.

[Dar97] Darr, Ethics in Health Services Management, Third Edition, Health Professions Press, Inc. Baltimore, MD.

[Dem00] J.H. Dempster. Conflict Identification Thesis. M.S. Thesis, North Carolina State University, 2000.

[EB02] J.B. Earp and D.Baumer. Innovative Web Use to Learn about Consumer Behavior and Online Privacy. *Communications of the ACM*, forthcoming 2002.

[EP00] J.B. Earp and F. C. Payton. Dirty Laundry: Privacy Issues for IT Professionals, *IT Professional*, 2(2), pp. 51-54, March/April 2000.

[Epi00] *Pretty Poor Privacy: An Assessment of P3P and Internet Privacy* http://www.epic.org/reports/prettypoorprivacy.html, Electronic Privacy Information Center, June 2000.

[FB91] W.J. Fabrycky and B.S. Blanchard. *Life Cycle Cost and Economic Analysis*, Prentice-Hall, 1991.

[FIP73] The Code of Fair Information Practices, U.S. Dep't. of Health, Education and Welfare, Secretary's Advisory Committee on Automated Personal Data Systems, Records, Computers, and the Rights of Citizens, viii, http://www.epic.org/privacy/consumer/code_fair_info.html, 1973.

[FTC98] *Privacy Online: A Report to Congress*, http://www.ftc.gov/reports/privacy3/, Federal Trade Commission, June 1998.

[FTC00] Privacy Online: Fair Information Practices in the Electronic Marketplace. A Report to Congress. Federal Trade Commission, 2000.

[GHS00] J. Goldman, Z. Hudson and R.M. Smith. Privacy Report on the Privacy Policies and Practices of Health Web Sites, Sponsored by the California HealthCare Foundation, Jan. 2000.

[JBC98] Jarke, M., X.T. Bui and J.M. Carroll. Scenario Management: An Interdisciplinary Approach *Requirements Engineering Journal*, Springer Verlag, 3(3-4), pp. 154-173, 1998.

[Lam01] A. van Lamsweerde. Goal-Oriented Requirements Engineering: A Guided Tour, *IEEE 5th International Symposium on Requirements Engineering (RE'01),* Toronto, Canada, pp. 249-261, 27-31 August 2001.

[PFI99] *Policy Framework for Interpreting Risk in eCommerce Security*. CERIAS Technical Report, http://www.cerias.purdue.edu/techreports/public/PFIRES.pdf, Purdue University, 1999.

[Pot99] C. Potts. ScenIC: A Strategy for Inquiry-Driven Requirements Determination, *Proceedings IEEE 4th International Symposium on Requirements Engineering (RE'99)*, Limerick, Ireland, 7-11 June 1999.

[RBP91] J. Rumbaugh, M. Blaha, W. Premerlani. F. Eddy and W. Lorensen. *Object-Modeling and Design,* Prentice Hall, New York, NY, 1991.

[RC97] J. Reagle and L. F. Cranor. The platform for Privacy Preferences. *Communications of the ACM*. 42(2), pp.48-55, February 1997.

[Rob97] W.N. Robinson. Electronic brokering for assisted contracting of software applets, *Proceedings of the Thirtieth Hawaii International Conference on System Sciences*, Vol. 4 , pp. 449-458, 1997.

[RSB98] C. Rolland, C. Souveyet and C.B. Achour. Guiding Goal Modeling Using Scenarios, *IEEE Transactions on Software Engineering*, 24(12), pp. 1055-1071, December 1998.

[Wil01] C. Wilson. "Lilly reveals Prozac patients' identities." http://www.infobeat.com/cgi-bin/WebObjects/IBFrontEnd.woa/wa/fullStory?article=409190643, 17 July 2001.