"I Need It Now": Improving Website Usability by Contextualizing Privacy Policies

Davide Bolchini[1], Qingfeng He[2], Annie I. Antón[2], and William Stufflebeam[2]

[1] TEC lab - Faculty of Communication Sciences - University of Lugano, Switzerland
davide.bolchini@lu.unisi.ch
[2] Requirements Engineering Research Group, NC State University, North Carolina, USA
{qhe2, aianton, whstuffl}@ncsu.edu

Abstract. Internet privacy policies are complex and difficult to use. In the eyes of end-users, website policies appear to be monolithic blocks of poorly structured texts that are difficult to parse when attempting to retrieve specific information. In an increasingly privacy-aware society, end-users must be able to easily access privacy policies while navigating a website's pages and readily understand the relevant parts of the policy. We propose a structured methodology to improve web design and increase user's privacy awareness. This systematic approach allows policy makers to effectively and efficiently reshape their current policies by structuring policies according to the subject that is relevant to specific user interaction contexts, making them more user-centered and user-friendly. The methodology is built upon prior work in privacy policy analysis and navigation context design.

1 Introduction

Privacy has become a more and more important issue and has recently received a lot of attention from consumers, government officials, legislators, and software developers due to concerns about increasing personal information collection from customers, information disclosure to third parties without user consent, and information transfer within and across organizations [5, 7, 8, 9].

Nowadays, most companies and organizations have posted one or more privacy policy documents on their websites. A privacy policy is a comprehensive description of a website's practices on collecting, using and protecting customer information. A privacy policy should define what information is collected and for what purpose, how this information will be handled, stored and used, whether customers are allowed to access their information collected by the website, how to resolve privacy-related disputes with this website, etc [6].

Unfortunately, current privacy policies published on websites are usually long and increasingly complex and difficult for users to understand. Research has found that many online privacy policies lack clarity and most requires a reading skill considerably higher than the Internet population's average literacy level [1]. There is a need to improve the current web design to help Internet users better navigate and understand website privacy policies and increase users' privacy awareness.

N. Koch, P. Fraternali, and M. Wirsing (Eds.): ICWE 2004, LNCS 3140, pp. 31–44, 2004.

We further elaborate this problem from the following four aspects of privacy policies: (1) content, (2) structure, (3) navigation, and (4) accessibility.

1. *Content*. The language used in privacy policies is often difficult for users to understand (e.g. either too technical or too legal), thus preventing them from easily understanding the benefits and potential threats entailed by the submission of their personal data. As recent studies have demonstrated [1], website privacy policies are often ambiguous and conflicting, and therefore preventing users from understanding how their personal information will actually be treated.
2. *Structure*. Different websites use different ways to present their privacy practices to users. For example, some policies (see www.bn.com) firstly explain *what* information they collect, and then *how* the organization will use and share this information. Other policies (see www.buy.com) tell *where* on the site they will collect user information and then focus on the strategy and technology used to *protect* that information. Other sites (see www.amazon.com) organize their policy's content with a list of frequently asked questions (FAQs), abruptly varying from very general issues (such as the kind of information collected) to technical details (e.g. the use of cookies) in the attempt to promptly answer the recurrent issues raised by the website's customers. In most of the cases, whichever strategy is chosen to organize the content of the privacy policy, the structure presented to users takes the shape of a long document with several sections (sometimes split into different physical web pages). Putting all of a privacy policy's information into one document may be useful to get a general overview of the site's privacy practices, however, having such a structure, policy texts are generally difficult to be contextualized into usage scenarios (e.g. inserting credit card information while buying a product) in which users may be concerned about the treatment of specific data (e.g. protection and storage of credit card number).
3. *Navigation*. With a monolithic structure such as this, privacy policy *navigation* is context-independent: wherever a user is navigating on the site, she can only access the entire privacy policy document as it is. No matter what the user is doing on the site, the policy always tells the same story in the same order. The question is, is that really what the user needs? For example, if a user connects to a site and realizes that she is promptly recognized personally by the site as a returning customer (e.g. "Hi <user's name>, here are our recommendations for you), she may wonder how (and for how long) her session data and shopping habits are stored and used by the organization. To reach such information, the user must go to another page and read a long, and possibly confusing document explaining in very general terms the importance of privacy, the effort spent by the organization to protect personal data, the technology used, and the conditions of use of her personal data (any of which may or may not be relevant for the described context of use), etc. Because of this, it is clear that users in a situation like this are presented with significant hurdles to retrieve the information they are interested in, and thus will more likely make the decision to blindly proceed with their site visit, being uninformed as to how their personal information will be used.

4. *Accessibility*. *The accessibility* of privacy policies is also usually very poor. The link to the privacy document is often difficult to spot, many times being designed as a recurrent pattern, in small font at the very bottom of the page. Even if it is accessible from every page of the site, it is not relevant to specific website pages. Once accessed, the privacy policy is still difficult to parse when attempting to retrieve specific information, as discussed in the previous bullets.

Taking these four dimensions of the problem into account, we can now formulate more clearly the specific problem we wish to address in this work.

In general, concerning the requirements for a "usable" online privacy policy, we argue that *users should be assisted while browsing or shopping on the site by understanding the privacy issues relevant to the current context of interaction*. More specifically, we should *provide users with **direct** access to **relevant** portions of the privacy policy concerning the information exchanged within the current context*. We propose that instead of "tell me the whole story about this organization's privacy practices", websites policies should "tell me ***now*** how the site treats my data that I'm ***now*** concerned with".

This paper proposes a systematic approach, which is built upon an existing goal-based policy analysis method, to address the aforementioned issues. Our approach allows policy makers and web designers to reshape their current privacy policies according to subject matters, thus meeting the specific needs relevant to the contexts of users' interactions. The expected benefits of applying this structured methodology for policy design are in two aspects. On one hand, users can have a better understanding of websites' privacy practices by accessing the relevant information quickly and increase their privacy-awareness. On the other hand, websites can build the trust from end-users by specifying privacy policies in an easy to access, easy to understand way to satisfy users' specific privacy needs.

The remainder of the paper is structured as follows: Section 2 discusses relevant previous work on goal mining, a powerful technique to examine and analyze privacy policy content. Section 3 details the proposed methodology and the expected results of the approach. Examples of application and concrete results are presented in section 4. Finally, section 5 summarizes the method, discusses some limitations of this approach and outlines our plans for future work.

2 Related Work

Privacy policy analysis has not been paid enough attention until recently. Studies following a structured approach to privacy policy analysis led to the development of specific analysis techniques based on goal-oriented requirements engineering practices [1, 2, 3, 4]. These conceptual methods and tools, which are based on goal-mining, turned out to be particularly effective for examining website privacy policies.

Goals are objectives and targets of achievement for a system in requirements engineering. In the case of privacy policies, a goal describes a statement expressing a privacy practice having a coherent and unitary meaning. Goal mining refers to extracting goals from data sources (in this case, privacy policies) by applying goal-based requirements analysis methods. The extracted goals are expressed in structured

natural language in the form of "VERB object" such as "COLLECT site usage information", "PREVENT storing credit card information using cookies", etc.

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?"*

All action words are possible candidates for goals. Goals in privacy policies are thus also identified by looking for useful keywords (verbs). 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 "Bank of America" Privacy Policy (www.bankofamerica.com):

> *Privacy Policy #1*: Employees are authorized to access customer information only when they need it, to provide you with accounts and services or to maintain your accounts.

By asking the goal identification questions, we identify the goal **G144**: *PROVIDE access to CI (Customer Information) to authorized personnel with authorized roles* from Privacy Policy #1.

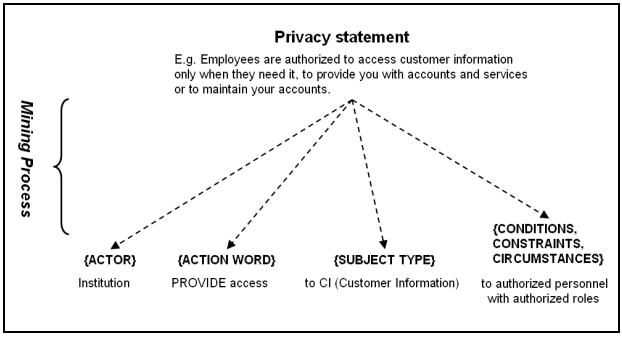

Fig. 1. Goal mining.

Figure 1 shows how a privacy statement is decomposed into four basic components of a privacy goal during the goal-mining process: *actor*, *action word*, *subject type*, and *conditions-constraints-circumstances*. The actor represents the stakeholder responsible for the goal to be achieved; the action word represents the type of activity described by the statement; the subject type describes the kind of user information at issue; finally, a goal usually recounts the conditions under which the goal actually takes place, the constraints to be respected, or other circumstances that provide the context to establish the scope of the goal.

The goal-mining process and the subject classification serve as the basis of the method proposed in this paper. This process of discovery, decomposition and representation cannot be entirely automated because it requires significant semantic content analysis. Goal-based analysis is best carried out by a team of analysts who do not simply chop each statement in a policy into pieces, but who carefully extract each statement's meaning, thus specifying goals that truly reflect the meaning of the

original document. However, tool support can greatly enhance the efficiency of the goal mining process[1].

3 Crafting Usable Policies

Based on previous discussion, we believe that users may benefit from directly accessing the corresponding parts of the privacy policy that are relevant to the web page a user is currently on. In this section, we propose a method to structure website privacy policies according to the subject matters and make them easily accessible to users. The goal of the proposed methodology is to make the transition from monolithic, poorly structured privacy documents to agile units of privacy policy content relevant to the current usage scenarios.

3.1 A Process Overview

We have identified five steps that will help designers create context-dependent policies (see Figure 2 for an overview):

1. Analyze existing privacy policy and identify privacy goals (goal mining)
Goal mining is the first step of this process and it enables analysts to gather a repository of privacy goals that represents the organization's privacy policies. This process is shown as step 1 in Figure 2. The process, techniques and heuristics to extract privacy goals from policy statements were detailed in [2, 3, 4].

2. Organize goals by subject type
These goals may be organized according to different criteria. For example, goals may be clustered by actor, action word and subject type. Organizing goals by subject type, which describes the kind of user information that a goal concerns, appears to be a very promising strategy for our purposes. Examples of subject types are, for example, *PII (Personal Identifiable Information)*, *credit card information*, *session data*, *purchase history, shipping data*, *account data*, *user preferences*, *usage habits* and *shopping habits*, *authentication information (e.g. user name and password),* etc. Other more domain-dependent subject types may be explored by analyzing policies from different application domains and business. For banking websites, for example, "bank account information" is particularly relevant, whereas it would not be relevant for B2C (business-to-consumer) type of e-commerce websites. In this step, we structure the collection of goals produced in step 1 according to the subject type. Each subject type is associated with a set of goals. It is noted that a goal could belong to more than one subject type. This process is shown as step 2 in Figure 2.

[1] In our approach, extracted goals are then documented in our Privacy Goal Management Tool (PGMT), a web-based tool developed at North Carolina State University. PGMT maintains a goal repository for analyses of policies and other documents from which goals can be derived. Each goal is associated with a unique ID, a description, a responsible actor, its sources and a privacy taxonomy classification.

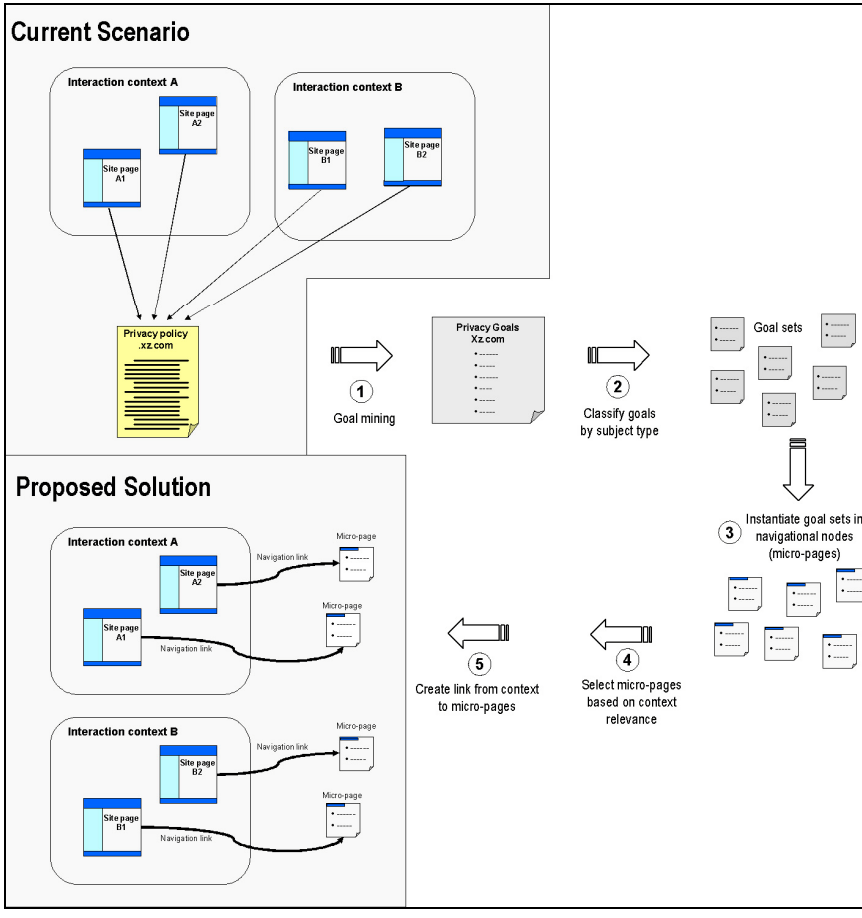

Fig. 2. A process overview for contextualizing privacy policies

3. Create a node for each goal set

Each set of goals (having the same subject type) may be compiled into a navigational node, a micro policy web page recounting the privacy goals in natural language.

4. Identify one or more contexts of user interaction that are relevant to a subject classification and associate them with the navigational node concerning that subject matter.

This is a crucial step, in which designers and policy makers must put themselves in user's shoes and envision the usage scenarios in which a user may need information about the organization's privacy policies (e.g., opening an account, purchasing a product, accessing personal information, modifying personal profile, etc.). Two lines of inquiry that may lead this task are:

a) *"For which task will a user need privacy policy information?"* The scenarios identified by answering this question usually take place in a

given context of interaction, i.e., while a user is browsing a certain set of web pages. Therefore, the second line of inquiry becomes relevant.

b) *"Where in the site will a user need privacy policy information to accomplish his/her task"?* For example, on the page where a user is entering credit card information or on the set of pages concerning the selection of shipping and payment preferences.

Thus, the result of this step is the identification of contexts of interaction where users may need privacy information (the *when* and the *where*).

Of course, each context of interaction should be associated with the privacy policy content relevant for users to accomplish the task in a given context, and since we have created a navigational node for each potentially relevant subject type (see step 3), we can associate one or more goal sets with each navigational context identified.

5. Create a link from each page of a navigational context to the associated goal set(s).

Once an interaction context has its goal sets associated with it, it is necessary to create links from the pages of the interaction context to the pages of the associated goal sets. Once this is done, a user may easily and directly navigate from a given page to the policy information relevant to the task she is doing. The relationship between goal sets and context of interaction is bidirectional. On one hand, an interaction context (e.g. shopping cart) may be associated with multiple goal sets (e.g.- the privacy goals concerning "buying history" and the goals concerning "user buying preferences"). On the other hand, the same goal set (e.g. goals concerning "buying history") may concern several interaction contexts, such as "access to homepage" (where recommendations are provided on the basis of user's buying habits), "shopping cart" (where related items are provided), "wish list", "customized pages", etc.

Links to privacy policy micro-pages should be clearly visible and easily accessible to users while they are performing a task (i.e., not at the very bottom of the page in small font).

3.2 Modeling Expected Results

We now discuss the expected outcomes of this process. Consider a navigation context such as the "Purchase process" in a generic e-commerce website. After having selected one or more products to buy, user is typically guided through a number of steps to complete the transaction. Each step is usually setup in a navigational node (an individual web page in this case). In some of these nodes, the user is asked to enter, confirm or modify the information concerning the transaction: in one page, the user has to enter payment information (credit card number, expiration date, type of card, name of the cardholder); in the subsequent page the user has to enter shipping information such as full address for delivery, delivery methods, and so on.

Currently, if the user wants to know more about the collection, storage and the security of payment information when he is on the "payment information" page, he has to scroll down to the bottom of the page, spot the small "privacy notice" link and start reading the long policy document while trying to find some clues and keywords to reveal the content of interest.

To overcome this problem, in our approach, a clearly visible link (for example with a text anchor like "privacy of payment information") is placed on the "payment information" page and leads the user directly to the relevant privacy micro-page (e.g. a side page unit or a pop-up) telling the user how payment information is handled and stored (see Figure 3).

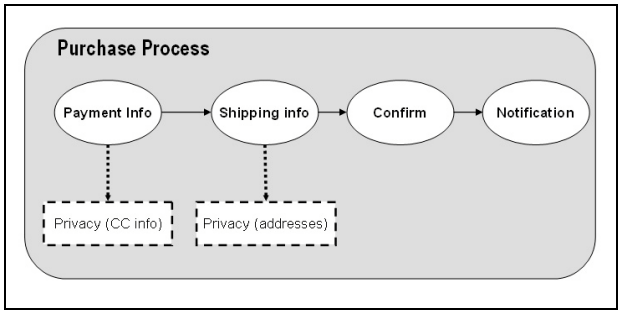

Fig. 3. Contextualized policy for the "Purchase process".

Additionally, in our approach policy may not only be made more accessible to the user, but the site might also raise awareness in the user about less-evident privacy practices, such as organization's privacy practices on session data. It is the case, as mentioned before, that the data about the user sessions (such as session time, IP, type of browser, information stored in cookies) are often associated to PII (Personal Identifiable Information) such as name, email, address, etc. Providing direct access to privacy information about these kinds of data helps raise user's privacy awareness on protecting their PII.

Likewise, on the homepage (which is often personalized through the use of session data) it may be relevant to have links to policy micro-pages concerning the treatment of session data. This may help users understand the reasons why the site gets increasingly customized as users access and provide information on various pages of the site, and for what other purposes this information is used by the organization.

The previously presented scenarios are intentionally generic to highlight the wide scope of applicability of the proposed methodology. The next section will focus on application examples taken from a well-known e-commerce website, thus defining more specific solutions and emphasizing the benefits for the user experience.

4 Application Examples

To demonstrate our approach, we now present some examples taken from an analysis of Amazon.com [10]. We have chosen this application because Amazon.com is a successful and typical e-commerce website familiar to most Internet users, which has a quite complex privacy policy and which may really benefit from adopting a contextualized perspective on its privacy communication. Although this is a specific case, most of the situations presented are likely common to many websites that gather user data for online transactions and better communication with their customers.

For each relevant interaction context, we will present contextualization solutions by detailing the following aspects:

- **Interaction context:** *web pages where the user may need specific and relevant policy information*
- **User issues:** *possible concerns of the user while navigating in the interaction context*
- **Link to relevant policy:** *a link available to the user pointing to relevant privacy policy micro-pages*
- **Policy micro-page content:** *the specific privacy information contained in the linked micro-pages (in terms of relevant privacy goals)*

The discussion of the examples is based on the assumption that web designers have applied our methodology described in Section 3 and produced sets of privacy goals according to the subject matters.

Scenario 1:

Let us consider the scenario in which a first time user of Amazon (a non-customer) connects to the site and declares her interest in a product by putting it in the shopping cart and then proceeds to check out. At this point, Amazon asks the user whether she is already an Amazon customer or not. The registration page for a new customer is shown in Figure 4.

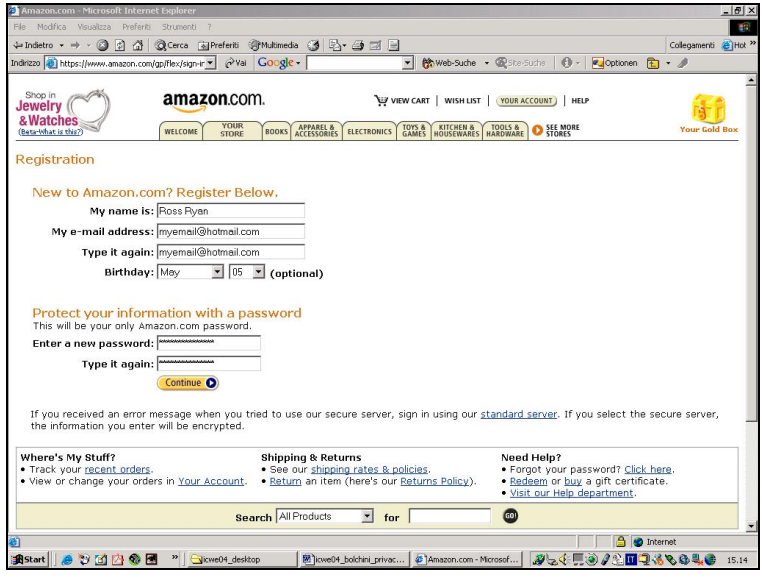

Fig. 4. Registration page for a new customer.

On this web page, a privacy-aware user might ask this question: how will Amazon use my name, email, date of birth and password? Currently, to clarify her issues the user has to know that at the very bottom of the page there is a link "Privacy Notice" that opens a page starting with the following section "What Personal Information About Customers Does Amazon.com Gather?" and then goes on with "What About

Cookies?" First, it is unlikely that a first time user (especially a non-frequent web surfer) knows a priori that by scrolling to the bottom of the page she will find the privacy policy link. Secondly, the privacy policy, as it is, presents the user with information that is completely irrelevant to her current context of use: the user already knows what kind of information Amazon is collecting from her; moreover, details on cookie use are not of interest to the user here. In the other two scenarios described in this section, there exists similar situation where the user has to scroll down to the very bottom of this page, find the small "Privacy Notice" link, and then read a long privacy policy statement to find the relevant information she needs to know.

Now let us examine this situation using our approach.

Interaction context (Figure 4): within the registration process, the user is on the new customer registration page.

User issues: "How will Amazon use my name, email, date of birth and password?"

Link to relevant policy: a link called "see how we treat your registration information" or "privacy for data exchanged in this form," positioned right beside or below the form.

Policy micro-page content: all privacy goals with subject matter {registration data}. Amazon's privacy policy contains the following privacy goals about this subject:

- G_{1349}: ALLOW customer to access personally identifiable information (including name, email, password, etc.)
- G_{1338}: AVOID companies and individuals who perform functions on our behalf using customer personal information for other purpose other than performing the specified functions
- G_{72}: AVOID selling customer information to others
- G_{748}: COLLECT information (e.g. personally identifiable information, assets, income, investment objectives, etc.) from forms submitted by customer (e.g. applications)
- G_{1339}: EMPLOY other companies and individuals to perform functions (such as processing credit card payments) on our behalf using customer information
- G_{88}: GUARD data during transmission using SSL encryption technology
- G_{1350}: SEND customer offers on behalf of other business without giving them name and address
- G_{639}: SHARE customer information among subsidiaries
- G_{168}: SHARE customer information related to your transactions with corresponding affiliates
- G_{492}: SHARE customer information as permitted by law
- G_{1132}: SHARE customer information with other organizations with customer consent
- G_{1351}: TRANSFER customer information as assets in case of buying/acquiring other companies or being acquired

Scenario 2:

In a different scenario, an existing customer may put products in the "Shopping Cart" while browsing the product catalog.

As soon as a product is put in the Shopping Cart, the user is presented with a page displaying suggestions of other potentially interesting products to put in the cart (Figure 5). Let us examine this situation.

Interaction context (Figure 5): Shopping cart suggestions.

User issues: "Amazon suggests several additional items for me to consider according to the purchase habits of other customers. For what other purposes will Amazon use information about my purchase habits?"

Link to relevant policy: a link called "privacy of your purchase history", positioned right below the page title "Customers who shopped for…also shopped for…"

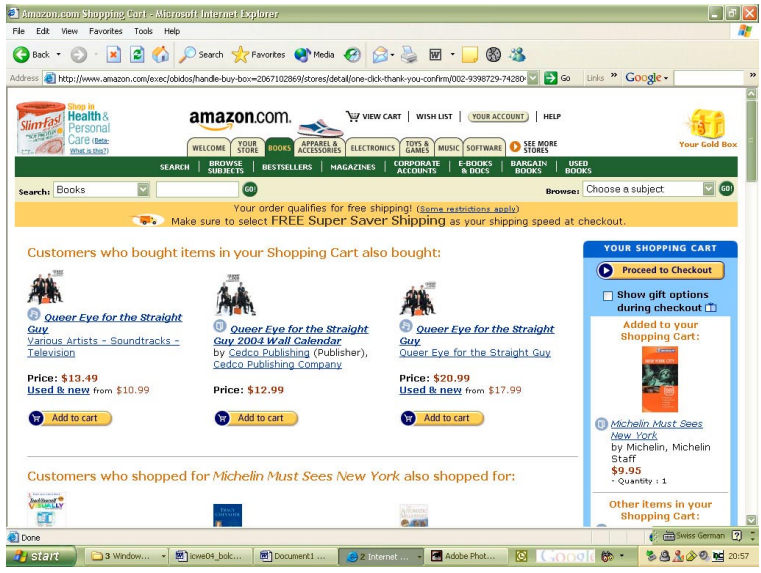

Fig. 5. Amazon Shopping Cart suggestion page.

Policy micro-page content: all privacy goals having as subject matter: {user history or previous purchases, etc}.

Amazon's privacy policy contains five goals about this subject:

- G1348: ALLOW customer to access recent product view history
- G1347: ALLOW customer to access recent purchase history
- G1344: ANALYZE purchase history
- G487: COLLECT information about customer online account (e.g. balances, transactions, email, bills, payment history)
- G1346: USE cookies to store items in your shopping cart between visits

Scenario 3:

Finally, let us consider one of the most debated issues in the treatment of consumer privacy: the handling and use of credit card information. Figure 6 shows the page where the user has to enter payment information such as credit card details.

Interaction context (Figure 6): "Shipping & Payment" page of the purchase process.

User issue: "How will Amazon use and protect my credit card information?"

Link to relevant policy: a link called "Privacy of credit card information" positioned right beside the question "Paying with a credit card?" (see Figure 8).

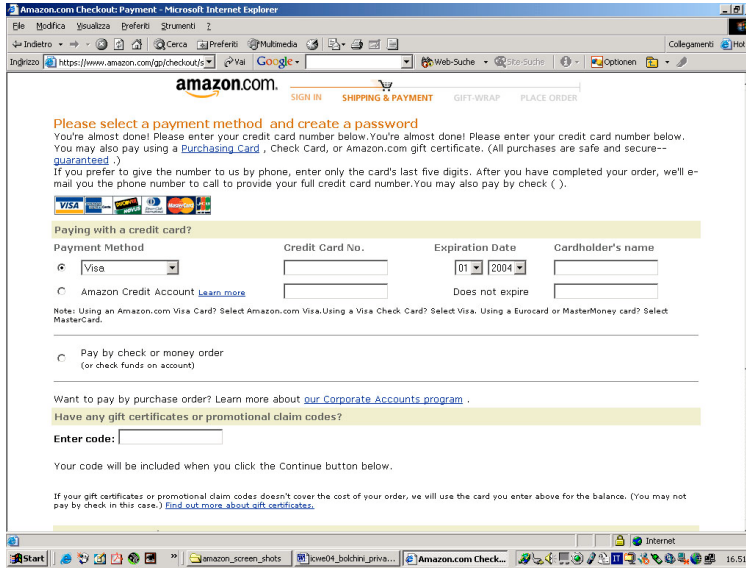

Fig. 6. Exchange of sensible payment information.

Policy micro-page content: all privacy goals having as subject matter: {credit card or user payment information}

Amazon's privacy policy contains the following privacy goals about this subject:

- G1337: ALLOW customer to access payment settings (including credit card information, etc.)
- G1338: AVOID companies and individuals who perform functions on our behalf using customer personal information for other purpose other than performing the specified functions
- G37: COLLECT credit card information for billing
- G1339: EMPLOY other companies and individuals to perform functions (such as processing credit card payments) on our behalf using customer information
- G88: GUARD data during transmission using SSL encryption technology

- G_{1341}: REVEAL only the last five digits of your credit card numbers when confirming an order
- G_{1342}: USE credit history information from credit bureaus to help prevent and detect fraud
- G_{1343}: USE credit history information from credit bureaus to offer credit or financial services to customers

Once selected, goals may be properly rephrased from their formal structure to a fluent narrative to make them more understandable for the user.

5 Summary and Future Work

Most online privacy policies are poorly structured, hard to understand, long documents that do not satisfy end-user's need for a concise policy statement for specific context. In this paper, we present a new method to analyze privacy policies to produce privacy goals, and then structure these goals according to the subject matters. By doing this, web designers can associate the context of web page to appropriate privacy goal sets that concerns the current subject.

This method is based on a structured and validated policy analysis process. This ensures the completeness and consistency of the goal sets displayed in policy micro-pages.

Both users and organizations can benefit from applying the proposed methodology in web design.

For users, they quickly get direct access to the relevant policy information at the right time (i.e. when they need it). This enhanced accessibility makes policies more and more visible to the users, thus raising overall awareness of Internet users to many privacy concerns. It also helps users evaluate the privacy practices declared by the organization in a much more straightforward manner.

Applying the proposed methodology (or even simply adopting the general idea), organizations can more easily evaluate the coverage of their privacy policy. By analyzing the different interaction contexts, site stakeholders have the opportunity to verify whether or not their policy contains information relevant for the user in that given context, not just generic and essentially useless information about privacy. A contextualized policy also builds trust of users to websites, since it communicates more clearly with site privacy practices, showing attention to the concrete needs of the user. Finally, contextualizing policies means enhancing the user experience on the site, providing more (or less) reasons for visitors to become customers.

The methodology we proposed and results gathered so far are even more useful for multi-channel applications, which are increasingly available on a variety of smaller devices, such as PDAs, handhelds, and smart phones. The visualization and interaction requirements of such devices pose more strict constraints to user's capability of interacting with and reading long documents such as privacy policy. In these cases, the contextualization and design of agile mini-policies are very important to make privacy policy really usable to end-users.

The method has yet to be empirically validated on large scale and across different domains through the validation of prototypes and usability testing. Moreover, return-

of-investment (ROI) for this change in policy communication still needs to be evaluated.

For future work, we are looking for other relevant interaction contexts where contextualized privacy policies can play a key role in improving user experiences. We are going to apply this methodology to other domains, such as banking and financial institutions websites, which we have already conducted goal-mining and privacy policy analysis studies.

One further extension of the approach that may be explored is that semantic associations between goal sets may further enhance privacy policy usability. For example, the "session data privacy" micro-page may be linked to the "recommendation system" privacy micro-page used by the organization, which exploits session data. Similarly, the "Recommendation system privacy" micro-page may be linked to the "shopping habits privacy" micro-page, since recommendations are built on previous shopping habits of the user, and so on. Such design solutions may lead to a privacy policy whose agile navigation highlights even more the semantics underlying the privacy practices.

References

[1] A.I. Antón, J.B. Earp, D. Bolchini, Q. He, C. Jensen and W. Stufflebeam. The Lack of Clarity in Financial Privacy Policies and the Need for Standardization, Accepted, to appear in: *IEEE Security & Privacy*, 2004.

[2] A.I. Antón and J.B. Earp. A Requirements Taxonomy to Reduce Website Privacy Vulnerabilities. To Appear: Requirements Engineering Journal, Springer Verlag, 2003.

[3] A.I. Antón, Q. He, and D. Bolchini. The Use of Goals to Extract Privacy and Security Requirements from Policy Statements, Submitted to: *the 12th IEEE International Requirements Engineering Conference (RE'04)*, January 2004.

[4] A.I. Antón, J.B. Earp and A. Reese. Analyzing Web Site Privacy Requirements Using a Privacy Goal Taxonomy. 10th Anniversary IEEE Joint Requirements Engineering Conference, Essen, Germany, 9-13 September 2002.

[5] J.B. Earp and D. Baumer. Innovative Web Use to Learn about Consumer Behavior and Online Privacy. *Communications of the ACM*, 46(4), April 2003.

[6] The Code of Fair Information Practices, U.S. Department 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.

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

[8] National Telecommunications and Information Administration. A Nation Online: How Americans Are Expanding Their Use of the Internet, http://www.ntia.doc.gov/ntiahome/dn/ Washington, D.C. February 2002.

[9] W.F. Adkinson, J.A. Eisenach and T.M. Lenard. Privacy online: A Report on the Information Practices and Policies of Commercial Web Sites. Washington, DC: Progress & Freedom Foundation, 2002. Downloaded July 18, 2003: http://www.pff.org/publications/privacyonlinefinalael.pdf.

[10] Amazon.com, Inc., http://www.amazon.com, Last visit: February 17th 2004.