Designing Natural Language and Structured Entry Methods for Privacy Policy Authoring

John Karat[1], Clare-Marie Karat[1], Carolyn Brodie[1], and Jinjuan Feng[2]

[1] IBM T.J. Watson Research Center, Hawthorne, NY 10532, USA
{jkarat, ckarat, brodiec}@us.ibm.com
[2] University of Maryland Baltimore County, Baltimore, MD 21250, USA
jfeng2@umbc.edu

Abstract. As information technology continues to spread, we believe that there will be an increasing awareness of a fundamental need to seriously consider privacy concerns, and that doing so will require an understanding of policies that govern information use accompanied by development of technologies that can implement such policies. The research reported here describes our efforts to design a system which facilitates effective privacy policy authoring, implementation, and compliance monitoring. We employed a variety of user-centered design methods with 109 target users across the four steps of the research reported here. This case study highlights our work to iteratively design and validate a prototype with target users, and presents a laboratory evaluation aimed at providing early support for specific design decisions to meet the needs of providing flexible privacy enabling technologies. This paper highlights our work to include natural language and structured entry methods for policy authoring.

1 Introduction

The rapid advancement of the use of information technology in industry, government, and academia makes it much easier to collect, transfer, and store personal information (PI) around the world. This raises challenging questions and problems regarding the use and protection of PI [13]. Questions of who has what rights to information about us for what purposes become more important as we move toward a world in which it is technically possible to know just about anything about just about anyone. As stated by Adams and Sasse [2]: 'Most invasions of privacy are not intentional but due to designers' inability to anticipate how this data could be used, by whom, and how this might affect users.' Deciding how we are to design privacy considerations in technology for the future includes philosophical, legal, and practical dimensions – any or all of which can be considered as within the domain of the field of human-computer interaction (HCI).

Privacy can and does mean different things to different people. We are primarily focused on a view of privacy as the right of an individual to control personal information use rather than as the right to individual isolation [15, 16, 22]. Organizations commonly provide a description of what kind of information they will

M.F. Costabile and F. Paternò (Eds.): INTERACT 2005, LNCS 3585, pp. 671 – 684, 2005.

collect and how they will use it in privacy policies. In some areas (e.g., the collection and use of health care information in the US or movement of personal information across national boundaries in Europe) such policies can be required, though the content of the policy is not generally specified in legislation. While there has been considerable consensus around a set of high level privacy principles for information technology [16], we do not think it is likely that a single privacy policy can be created to address all information privacy needs. For example, there will likely be considerable differences in privacy legislation in different regions of the world [14]. Similarly, organizations in different fields (e.g., healthcare, banking, government) need to tailor policies to their domains and needs [6, 7]. While we will focus on privacy policy, we acknowledge that privacy is not entirely about "setting rules and enforcing them" [18]. To implement privacy within an organization, the coordination of people, business processes, and technology is required. Still we do believe that privacy policies are essential when interacting with technology and/or organizations in that they enable people to better understand just how the boundary between public and private information is impacted by technology [3].

It is interesting to note that while privacy policy is not new to most organizations, very little has been done to implement the policies through technology [21]. Usability has been identified as a major challenge to moving the results of security and privacy research to use in real systems [8]. One reason seems to be that there has been only limited research into how to make complex security and privacy functionality understandable to those who must use it.

Privacy policy enforcement remains largely a human process, and privacy policies which organizations present to customers are generally very vague (e.g., "We will only use your personal information for the efficient conduct of our business"). There are emerging standards for privacy policies on websites [9], but these address machine readable policy content without specifying how the policy might be created or implemented. The reality is that there is very little capability to have technology actually implement access and use limitations we might expect from a policy statement like "We will not share your information with a third party without your consent". Our research focus has been on how organizations could create a wide range of machine readable policies, and how technology might enable the policies to be enforced and audited for compliance. We have elected to focus on technology to enable usable privacy policy authoring and enforcement, rather than trying to directly address what privacy rights people should have [e.g., 23] or how to de-identify information stored in systems [e.g., 20]. This does not mean that we think these aspects of privacy are not important social issues. Rather it points to our belief that technology can enable flexible, reliable and accountable privacy policy (i.e., be privacy enabling) and not just be a force which reduces individual rights. We hope our work contributes positively to this goal.

1.1 Privacy Policy Structure

Research from the International Association of Privacy Professionals (IAPP) reports that 98% of companies have privacy policies. Often organizations have both internal policies, which state rules about information handling within an organization, and external policies which describe the policy in terms intended to inform the data

subjects about use of their information. We focus here on internal policies, largely because they describe actual data handling procedures in organizations. These policies have been found to have a fairly specific structure which describes who can use what information for what purposes [19]. First of all organizations generally have a number of internal privacy policies; some to address use of data about internal employees, and others to address use of data about individuals with which the organization interacts (e.g., customers, patients, clients). Any policy includes a number of rules governing the use of data-subject's information. The rules in a privacy policy include data user, data element, purpose, use, condition, and obligations [5]. The first four of these elements can be said to be required of any good policy rule, and the last two are optional. The data user who accesses the data may be acting in a particular role in regard to a purpose. For example, doctors may read protected health information for medical treatment and diagnosis. In many privacy policies and legislation, granting or denying access incurs an obligation on the data user to take additional actions. For example, a medical researcher may read protected health information for medical research if the patient has previously explicitly authorized release (i.e., the condition) and the patient is notified within 90 days of the release of information (i.e., the obligation).

1.2 Motivation for Our Privacy Research

Most organizations store PI in heterogeneous server system environments. Currently they do not have a unified way of defining or implementing privacy policies that encompass data collected and used by both Web and legacy applications across different server platforms [4]. This makes it difficult for the organizations to put in place proper management and control of PI, the data users to access and work with the PI inline with the privacy policies, and the data subjects to understand rights regarding use of their PI.

In this paper we present a case study of a user-centered design research program on organizational privacy capabilities. We employed a variety of usability methods to progress from identifying organizational privacy concerns and needs to designing and evaluating prototypes and design trade-offs. This work included four steps: (1) identifying privacy needs within organizations through email survey questionnaires, (2) refining the needs through in-depth interviews with privacy-responsible individuals in organizations, (3) designing and validating a prototype of a technology approach to meeting organizational privacy needs through onsite scenario-based walkthroughs with target users, and (4) collecting empirical data in a controlled usability laboratory test to understand the usability of privacy policy authoring methods included in our proposed design. These activities were completed between the spring of 2003 and summer of 2004 and involved participation of 109 target users from around the world. From the first two steps we identified organizational needs which guided us in our choice of a focus area for the design of a system to improve privacy management for organizations. We focus our presentation here on our prototype development to meet these needs along with a laboratory study to evaluate the feasibility of our direction.

2 Designing and Evaluating a Privacy Policy Prototype

We designed and developed a prototype of a privacy policy workbench called SPARCLE. The overall goal in designing SPARCLE was to provide organizations with tools to help them create understandable privacy policies, link their written privacy policies with the implementation of the policy across their IT configurations, and then help them to monitor the enforcement of the policy through internal audits. Once we designed a prototype, we conducted a series of walkthrough sessions in which we utilized the prototype to discuss an appropriate scenario with representatives of healthcare, government, and finance organizations. In this paper, we will concentrate on the techniques we designed and developed for authoring privacy policies and assisting organizations in understanding the policies that have been created. While we present work on authoring policies in English, the approach and underlying technology allows the development of similar systems for other languages.

2.1 Designing a Prototype for Authoring Privacy Policies

During the survey and interview research, many of the participants indicated that privacy policies in their organizations were created by committees made up of business process specialists, lawyers and security specialists as well as information technologists. Based on the range of skills generally possessed by people with these varied roles, we hypothesized that different methods of defining privacy policies would be desirable. Our design direction was to support users with a variety of skills by allowing individuals responsible for the creation of privacy policies to define the policies using natural language or to use a structured format to define the elements and rule relationships that will be directly used in the machine-readable policy. SPARCLE keeps the two formats synchronized. For users who prefer authoring with natural language, SPARCLE transforms the policy into a structured form so that the author can review it and then translates it into a machine-readable format such as EPAL [5]. SPARCLE translates the policies of organizational users who prefer to author rules using a structured format into both a natural language format and the machine-readable version. During the entire privacy policy authoring phase, users can switch between the natural language and structured views of the policy for viewing and editing purposes. Once the machine-readable policy is created, it is possible to create enforcement engines to ensure the policy is enforced for data stored in the organization's on-line data stores.

Figure 1 contains a screen capture of SPARCLE's natural language interface for deEfining privacy policies. Throughout SPARCLE, the tool provides a task flow in the form of tabs showing the high level task steps to be accomplished and the status of each. The tasks include: Author Policy (step shown in Figures 1 and 2), Transform Policy (step shown in Figure 3), Map User Categories, Map Data Categories, Map Purposes/Actions, Map Conditions, Map Obligations, and Verify Policy. The mapping steps are used to associate policy elements with system objects, and enable separation of high level and detailed policy specification. The page also contains general information about the policy, (the name, date created, and file source of the policy, and a description of the policy authoring task to be performed) a list of privacy

policy templates that could be either provided by the tool for particular domains and geographies based on laws or created by the organization for customization and use by its divisions, and an Example Rule Guide describing the elements that make up a privacy policy rule. The privacy policy rule guide is based on analyses of privacy policy rules specified in [5].

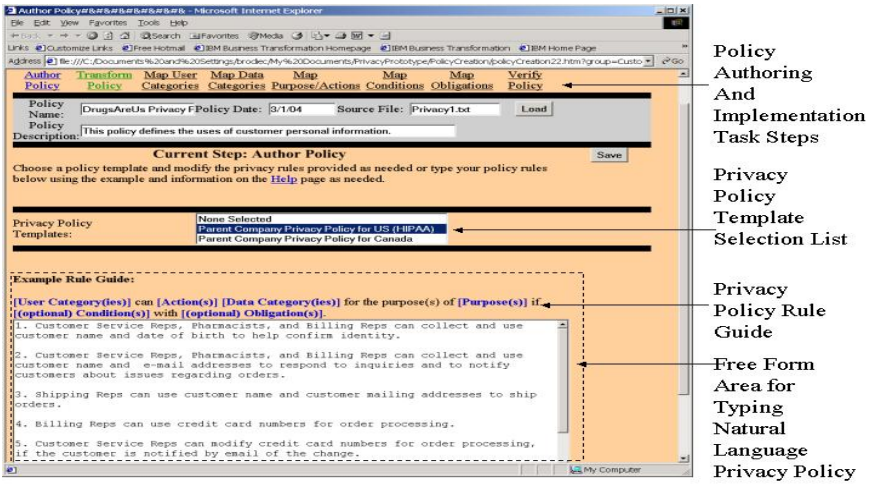

Fig. 1. SPARCLE natural language privacy policy creation screen

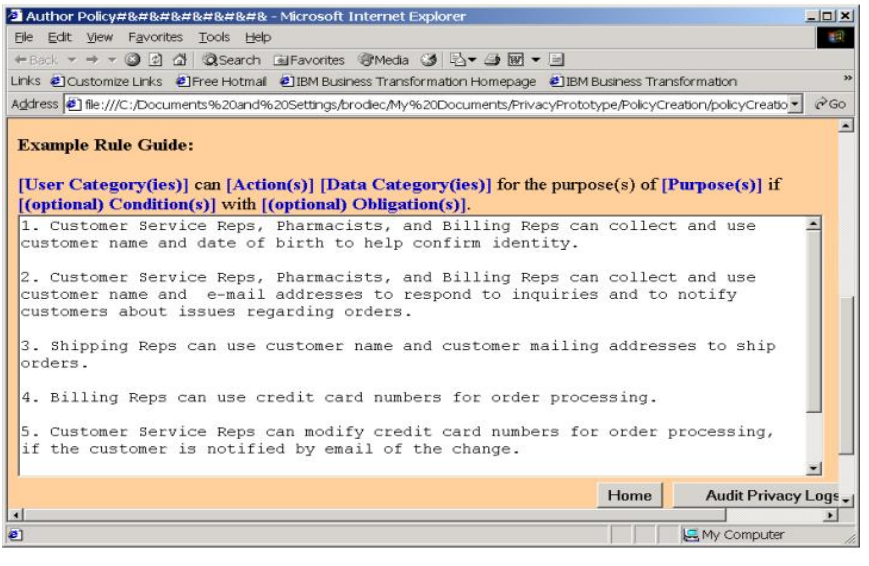

Fig. 2. Expanded view of natural language policy guide and entry field

The guide defines the basic components that are necessary in a privacy policy rule that is enforceable including user categories, allowed actions, data categories, purposes, as well as optional components such as conditions and obligations. Finally, a text entry area is provided for the actual privacy policy. When the user begins the process of creating a new policy, she can create the policy from scratch by typing into the text entry area, copy an existing policy into that area, or select one of the templates provided and modify it. The portion of Figure 1 within the dotted lines is enlarged and shown in Figure 2.

When the author is satisfied with the policy, he clicks on the Transform Policy tab shown in Figure 1. The natural language policy is analyzed and the policy elements (the strings which describe the User Category, Action, Data Category, Purpose, Conditions, and Obligations) in each rule are identified using a natural language parser (a shallow parser with a grammar and a domain dictionary). The natural language entry field area is replaced with a structured privacy policy creation view (shown in Figure 3). On this page, the user is provided with a list containing the parsed rules in the current policy.

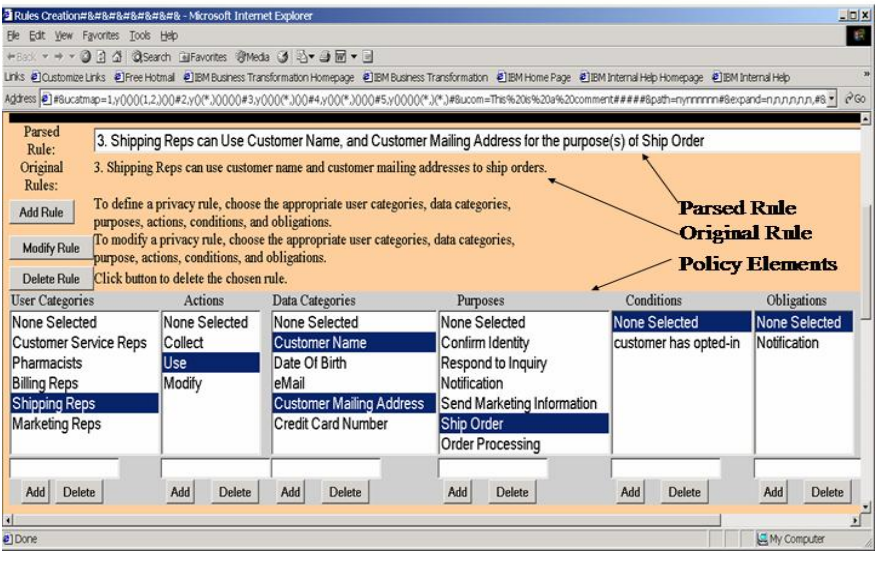

Fig. 3. Expanded view of SPARCLE structured privacy policy rule creation

Whenever a parsed rule is selected in the transformed view, the original unparsed text is also displayed and the elements of the rule that have been identified are highlighted in individual policy element selection lists as shown in Figure 3. There is one policy element selection list for each of the 6 types of rule element. There were two original purposes for this part of the prototype. First, while the natural language parsing technology in a limited domain such as privacy policy creation has promising accuracy, it is not perfect. This page allows users who have created the policy using the natural language technique to confirm that the parsing technology has identified

all parts of the rules correctly and to correct anything that is in error. Second, for users who prefer the more structured method for privacy policy creation, this method can be used to create the entire policy. The organization can define policy element lists and then rules can be created by selecting the appropriate elements from each of the policy element selection lists and selecting "Add Rule". Likewise, a rule can be modified or deleted by highlighting the rule in the rule selection list, modifying the selected elements as appropriate and selecting "Modify Rule" or "Delete Rule". Any modification to rules or rule added or deleted using the structured approach is automatically reflected in the natural language version of the rules as well. Therefore, the author is able to go back and forth between the two methods to view the policy either in natural language or the parsed format with the elements identified.

During the course of the scenario-based sessions with target users, they identified an additional use of the combined natural language and structured methods. The users indicated that the natural language parsing and display of parsed policies would be valuable to them for assessing the completeness of their existing privacy policies. Several participants were excited about the possibility of using SPARCLE to analyze their existing natural language privacy policies and then view the elements and rules identified in order to identify gaps and inconsistencies in the policies. For example, if an existing privacy policy rule fails to identify the purpose for which a particular user group is allowed to use a particular piece of data, the parsed rule would contain "none found" where purpose would usually be. The organizational users felt that this would be a valuable tool to ensure the quality of the privacy policies used by the organization and helpful in educating their organizations regarding their privacy policies.

Based on the data collected from interviews with organizational users responsible for the creation of privacy policies, they often find it difficult to understand the policies that they create in order to ensure that policies are complete, able to be implemented, and consistent. Figure 4 shows our design to provide users with easy ways of viewing the privacy policy. The figure contains a table in which two of the policy elements types are used as axes and the other elements that are associated with each row and column are shown in the cells. In the example that is shown, user categories are placed on the horizontal axis and data categories are placed on the vertical access. The cells in the table contain the purposes, conditions, and obligations for rules that apply to that user and data category. Using this table, users can see at a glance what type of users are allowed to access each data element and also see which user groups are never allowed to access particular data items. While the table format was well received by users, we are not yet sure how well a two dimensional table scales up to real organizational policy complexity. Scaling and visualization will be the subject of our future research.

2.2 Validation of Prototype with Target Users

We conducted scenario-based usability walkthrough sessions of two iterations of SPARCLE with people who were responsible for the creation, implementation, and auditing of privacy policies within large organizations in the domains of healthcare, banking, and government. During the course of the 90 minute sessions with 1 to 4 participants, we gathered verbal and written feedback on the usability, design, and

value of the privacy tool. For the first iteration of the prototype, walkthrough participants (7 participants in 5 sessions) rated the prototype positively (an average rating of 5.4 on a 7-point scale with 7 indicating "highest value" and 1 indicating "no value"). We present this summary result since it communicates the overall response to the prototype. However, the primary purpose for the sessions was to gather more qualitative responses from the participants about the value of the system to their task of managing privacy (some of which is described below).

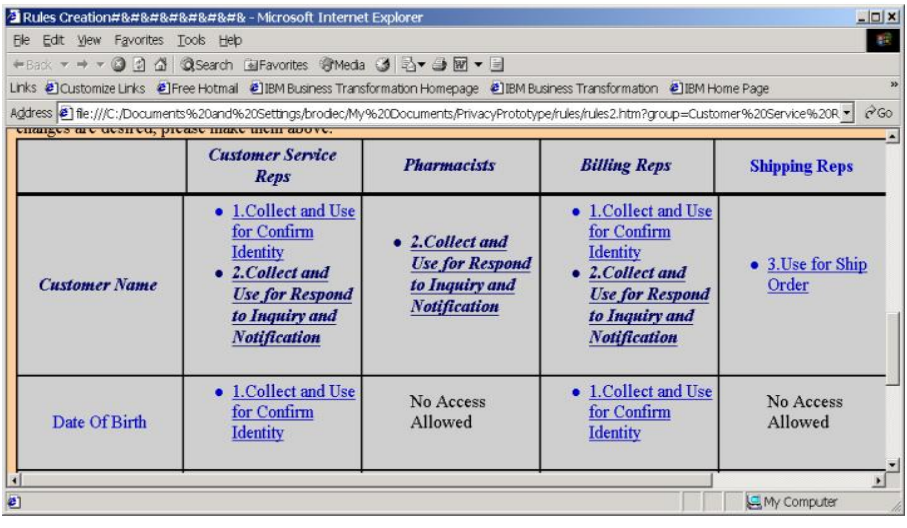

Fig. 4. Table showing privacy policy rules that apply to each user and data category

At the conclusion of the first iteration of design and evaluation, we made the following changes: 1) We added the ability to import pre-existing privacy policies into the natural language policy authoring condition to allow SPARCLE to highlight gaps and inconsistencies in the policies, 2) We added the ability to use privacy policy templates as a starting point for authoring privacy policies using either the natural language or structured policy authoring methods, and 3) We improved the readability of the table view of the privacy policy by bulleting entries and making it scrollable (See Figure 4). Additional improvements were made to the mapping and auditing functionality which we will not discuss here. Based on the feedback from our walkthrough sessions, we also decided to conduct an empirical test of the two authoring methods described in Step Four. During the second iteration of walkthrough sessions, the participants (15 participants in 6 sessions) also rated the revised prototype very positively (an average rating of 5.5 on the same scale).

The evaluations included 20 features on which we wanted to obtain feedback from the target users. Figure 5 summarizes the evaluation results over the two iterations of the prototype for 5 of these features which were included in both versions of the prototype and one feature that was added for the second iteration. While the data presented here only represent a small sample, we think that it provided us with a good

picture of how the users responded to the prototype. The added feature was the ability to import policy files from other sources and to modify those files. This would enable localization of larger corporate policies or laws. This was seen as a highly valuable feature in itself, and we also believe that it led to a more positive evaluation of the natural language entry in the second iteration of SPARCLE. While structured rule entry seemed to be preferred in the first iteration, Natural Language and Structured Entry had equal ratings in the second iteration (these features were not altered substantially between iterations). It was also important to hear from the target users that they felt there was considerable value in the fairly simply policy table that we included in the prototype. We had viewed this two-dimensional representation as an initial design which we might need to change substantially, but found that users actually found it to be very clear and a powerful tool for understanding policy coverage. Additionally, target users responded very positively to the incremental authoring process which allowed high level specification in natural language followed by detail specification (possibly by a different person at a different time). Finally, the target users provided feedback that the compliance checking capabilities we included in the prototype were in line with what they needed to offer end users details of how PI was being used within their organizations (by enabling records of accesses to specific user's information).

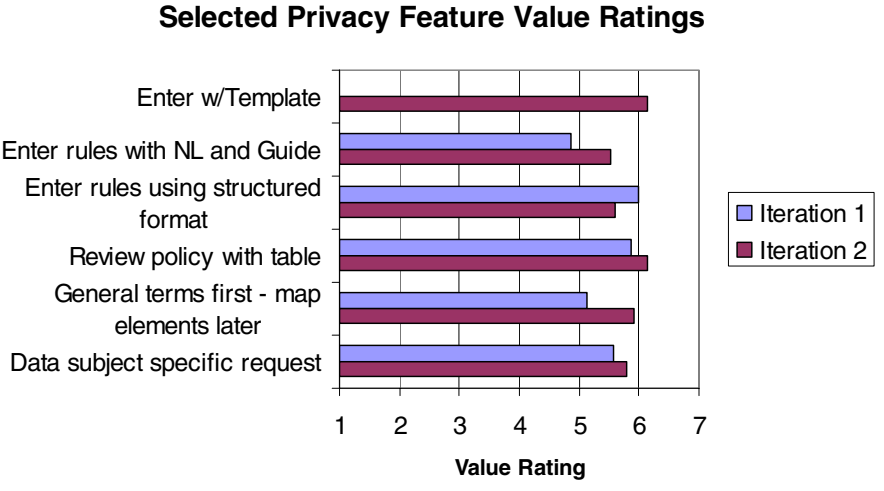

Fig. 5. Summary of Evaluation of Privacy Policy Authoring Features by Target Users

3 Evaluating Natural Language and Structured Policy Authoring

An empirical laboratory study was run to compare the two privacy policy authoring methods illustrated in the prototype. In order to provide a baseline comparison for the two methods (Natural Language with a Guide, and Structured Entry from Element Lists), we added a control condition that allowed users to enter privacy policies in text in any format that they were satisfied with (Unguided NL). The intention of the study

was not to strictly resolve which approach was better, but rather to inform the design by asking whether the two methods could easily be used to produce reasonable quality rules. While it is generally important to utilize target users in laboratory studies, we elected to use privacy policy novices in this study for two reasons. First, our earlier work with customers suggested that authoring privacy policies is undertaken by an audience with a variety of backgrounds and with no specific training in privacy policy authoring. We expect this to change over time, and that authors will become skilled as they gain access to the sort of tools we are developing. Second, the population of people skilled at writing policies which can be implemented is small. Thus we felt it practical and appropriate to look at the methods we were designing with a general audience. Our primary goal was to decide whether natural language authoring seemed promising enough to include in future research.

3.1 Experimental Design

Thirty-six employees of a large IT company were recruited through email to participate in the study. The participants had no previous experience in privacy policy authoring or implementation.

A repeated measures design was employed in the study; each participant completed one task in each of the three conditions. All participants started with a privacy rule task in the Unguided NL control condition (Unguided NL). Then, half of the participants completed a similar task in the Natural Language with a Policy Rule Guide condition (NL with Guide), followed by a third task in the Structured Entry from Element Lists condition (Structured List). The other half of the participants completed the Structured List condition followed by the NL with Guide condition.

In each task, we instructed participants to compose a number of privacy rules for a scenario we provided which described a desired privacy situation. Participants worked on three different scenarios in the three tasks. We developed the scenarios in the context of three privacy sensitive domains, namely health care, government, and banking. Each scenario included a description of a situation with five or six privacy rules (statements of who could use what information for what purpose), which included one condition (e.g., "If the customer agrees") and one obligation (e.g., "We will delete your personal information from our databases after one year"). The order of the scenarios was balanced across all participants.

We recorded the time that the participants took to complete each task and the privacy rules that participants composed. We also collected, through questionnaires, participants' perceived satisfaction with task completion time, quality of rules created, and overall experience after participants completed each task.

In order to compare the quality of the rules participants created under different conditions and scenarios, we developed a standard metric for scoring the rules. We counted each element of a rule as one point. Therefore, a basic rule of four compulsory elements had a score of four and a scenario that consisted of five rules, including one condition and one obligation, had a total score of 22. We counted the number of correct elements that participants specified in their rules, and divided that number by the total score of the specific scenario. This provides the percentage of elements correctly identified for comparison across different scenarios and conditions.

3.2 Results and Discussion

There was a significant difference in the task completion time across the three conditions ($F_{(2, 70)} = 4.58$, $p < 0.05$). Mean participant time on task was 910 seconds for Unguided NL, 814 secs. for NL with Guide, and 992 secs. for Structured List conditions respectively. Post hoc tests showed that the NL with Guide method took significantly shorter time than the Structured List method. There was no significant difference between the Unguided NL method and the other two methods.

A repeated measures test with post hoc analyses indicated that participants were more satisfied with the quality of the rules created by the NL with Guide method or the Structured List method as compared with the Unguided NL method ($F_{(2, 70)} = 6.54$, $p < 0.005$). On a scale of 1 to 7, with 7 indicating highest overall satisfaction, participants mean satisfaction scores were 4.0 for Unguided NL, 4.7 for NL with Guide and 4.6 for Structured List conditions. There was no significant difference between the NL with Guide method and the Structured List method.

A statistical test of the rule quality scores calculated using the standard metric found a significant difference between the three conditions ($F_{(2, 70)} = 44.3$ $p < 0.001$) (see Figure 6 below). Post hoc tests showed that the NL with Guide method and the Structured List method helped users create rules with significantly higher quality than the Unguided NL method. There was no significant difference between the NL with Guide method and the Structured List method. Using the Unguided NL method, participants correctly identified about 42% of the elements in the scenarios, while the NL with Guide method and the Structured List method users correctly identified 75% and 80% of the elements, respectively. Since we did not provide feedback on rule quality in any method, we attribute most of the improvement to the authoring methods themselves and not to learning in the first trail.

We examined the readability of the policies created. Jensen and Potts [11] found that privacy policies posted on the Web were generally not easy to read. We adopted the same measurement approach and used the Flesch readability score to evaluate the readability of the rules composed in the study. A repeated measures test suggested that there was a significant difference in the readability of the rules composed in the three conditions ($F_{(2, 70)} = 15.89$, $p < 0.001$). A post hoc test showed that the rules composed in the NL with Guide condition were significantly more difficult to read than the other two conditions. There was no significant difference between the readability of the rules created under the Unguided NL and Structured List conditions, and they were of similar readability as the majority of the online privacy policies reported by Jensen and Potts [11].

The results of the experiment confirmed for us that both the NL with Guide method and the Structured List method enabled participants to create rules with higher quality than the Unguided NL method. The fact that the percentages of elements identified with these methods almost doubled that of the Unguided NL method suggests that the NL with Guide and the Structured List methods are reasonably easy to learn and use. Our purpose in conducting the study was not simply to select one of the methods as the best method to include in the prototype. Certainly, if one of the methods seemed significantly superior, it would have made us consider going forward with a single method. However, our explorations of the two methods with customers had suggested that going forward with two methods that we complimentary might be a

preferred solution. We view the results of this study as giving support to the design direction of including both methods and allowing rule creation in either to accommodate author preferences.

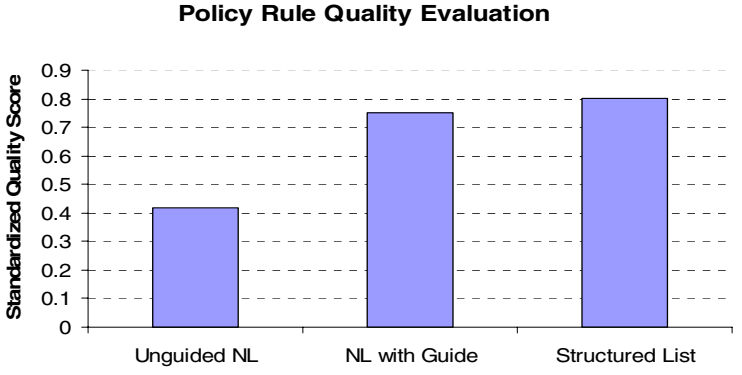

Fig. 6. Average scores of the quality of the rules according to the quality evaluation metric in three conditions

4 Conclusions and Future Research

Privacy is emerging as a powerful issue for people within organizations and individuals who interact with them around the world. In the networked world in which we live today, the topic is of growing concern and importance. Previous research has shown that the general public is concerned about protecting their privacy and often does not understand the implications of the privacy policies published by organizations with which they share their PI [11, 18]. This case study highlights the work of identifying organizational privacy requirements, iteratively designing and validating a prototype with target users in their work settings, and empirical laboratory testing to guide specific design decisions to meet the needs of providing flexible privacy technologies for organizations and their users.

Early work with privacy representatives in this project convinced us to focus on policy authoring, implementation and auditing in our research. We designed and developed a prototype with the overall goal of providing organizations a tool to help them create understandable privacy policies, link their written privacy policies with the implementation of the policy across their IT configurations, and then help them to monitor the enforcement of the policy through internal audits. We explored and iterated on the design with target users and were able to obtain valuable feedback well before we could complete a full implementation of the prototype. While work on the natural language parsing and mapping components of SPARCLE is still underway, we think we have a solid understanding of organizational needs for privacy technology.

We also conducted an empirical usability laboratory test of two methods of authoring policies. Results were promising and showed that in initial use, novice users could use the two methods to identify and cover 75-80% of the policy rule elements.

Coupled with our work with target organizational users, we have concluded that integrating the Structured List and NL with Guide authoring methods along with providing an easy to understand policy coverage view will be important elements of a successful privacy policy tool. We think that the laboratory test was an important component of the overall research in helping to justify the value of including a natural language method and integrating it with a structured authoring approach.

We think that a number of research challenges remain. First, we need to examine how well our authoring environment works for realistically complex organizational privacy policies. Our target users have generally been from large organizations, and they have responded well to the parts of the prototype we present in this paper – authoring and viewing policy coverage. However, working with policies with hundreds of rules might create problems that do not emerge in discussions centered on a single policy involving a few rules. A planned next step and a natural evolution for our work will be to work with several organizations to create complete machine readable policies which reflect their actual internal privacy policies. In doing this we hope to address issues about the use of internal policies in communicating with end users concerning privacy. We suspect that well formed internal policies will also be useful descriptions as external policy documentation. Related to this, is a belief that better tools for policy authoring can enable the creation of clearer privacy related legislation. We are still in a time where there is a considerable gap between what privacy laws say should be done and what technology actually can help make happen.

There are some challenges that future research and professional groups will need to address before our work could contribute to a generally useful privacy technology. First, standards need to advance beyond those currently in place [9] so that it becomes technically feasible for privacy policy information to travel with data within and outside of organizations. Perhaps a focus on the importance of privacy could contribute to changes in system architecture – to enable easier privacy and security. Current world events are providing pressure on the public and private sectors to consolidate data and collect and use more PI for a variety of purposes. At the same time, legislation in countries around the world is providing data users with increased obligations regarding the use of PI and data subjects with rights about the collection and use of their PI by organizations. Technology can help to protect people's privacy in collaboration with social policy. The HCI field can step up to the challenge of creating interfaces and interaction methods that reduce the complexity in defining, implementing, and managing privacy policies for the benefit of all parties.

References

1. Ackerman, M. Darrell, T., and Weitzner, D. (2001) Privacy in context, *Human Computer Interaction*, 16, 2, 167-176.
2. Adams, A. and Sasse, A. (2001) Privacy in Multimedia Communications: Protecting Users, Not Just Data . In A. Blandford, J. Vanderdonkt & P. Gray [Eds.]: People and Computers XV - Interaction without frontiers. *Joint Proceedings of HCI2001 and ICM2001*, Lille, Sept. 2001. pp. 49-64. Springer.
3. Altman, I. (1975). *The Environment and Social Behavior, Privacy, Personal Space, Territory and Crowding*. Monterey, CA: Brooks/Cole Pub. Co., Inc.

4. Anton, A., He, Q., and Baumer, D. (2004) The complexity underlying JetBlue's privacy policy violations. *IEEE Security & Privacy*. August/September, 2004.
5. Ashley, P., Hada, S., Karjoth, G., Powers, C., and Schunter, M. (2003). *Enterprise Privacy Architecture Language (EPAL 1.2).* W3C Member Submission 10-Nov-2003. http://www.w3.org/Submission/EPAL/
6. Ball, E. (2003). Patient privacy in electronic prescription transfer. *IEEE Security and Privacy*, 1, 2, 77-80.
7. Baumer, D., Earp, J.B., and Payton, F. C. (2000). Privacy in medical records: IT implications of HIPAA. *Computers and Society*, December, 2000, 40-47.
8. CRA Conference on "Grand Research Challenges in Information Security and Assurance". http://www.cra.org/Activities/grand.challenges/security/. November 16-19, 2003.
9. Cranor, L. (2002). *Web Privacy with P3P.* Cambridge: O'Reilly.
10. Hagen, P. (2000). Personalization versus privacy. *The Forrester Report*, Nov., 2000, 1-19.
11. Jensen, C. and Potts, C. (2004). Privacy polices as decision-making tools: An evaluation of online privacy notices. *CHI 2004*, 471-478.
12. Karat, C., Brodie, C., Karat, J., Vergo, J., and Alpert, S. (2003) Personalizing the user experience on ibm.com*. IBM Systems Journal*, 42, 4, 686-701.
13. Kobsa, A. Personalized hypermedia and international privacy. Communications of the ACM, 45, 5, 64-67.
14. Manny, C. H. (2003). European and American privacy: Commerce, rights, and justice. Computer Law and Security Report, 19, 1, 2003, 4-10.
15. National Research Council. (2003). Who goes there? Authentication through the lens of privacy. Washington, D.C: National Academies Press.
16. OECD (1980). OECD guidelines on the protection of privacy and transborder flows of personal data. http://www.oecd.org/home/
17. Office of the Federal Privacy Commissioner of Australia. (2000). *Privacy and Business (2000).* http://www.privacy.gov.au
18. Palen, L. and Dourish, P. (2002). Unpacking 'privacy' for a networked world, *CHI 2002*. 129-136.
19. Ponemon Institute and IAPP. (2004). 2003 benchmark study of corporate privacy practices.
20. Senior, A., Pankanti, S., Hampapur, A., Brown, L., Tian, Y., and Ekin, A. (2004). Blinkering Surveillance: Enabling Video Privacy through Computer Vision. *IEEE Security and Privacy*, in press.
21. Smith, J. (1993). Privacy policies and practices: Inside the organizational maze. Communications of the ACM, 36, 12, 105-122.
22. U.S. Fair and Accurate Credit Transaction Act. (2003). H.R. 2622, 108th Congress, July 24, 2003.
23. Warren, S.A. and Brandeis, L.D. (1890). The right to privacy. Harvard Business Review, Dec, 4, 195.