A User Study of the Expandable Grid Applied to P3P Privacy Policy Visualization

Robert W. Reeder
Microsoft
1 Microsoft Way
Redmond, WA 98052
roreeder@microsoft.com

Patrick Gage Kelley
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA, 15213
pkelley@cs.cmu.edu

Aleecia M. McDonald
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA, 15213
am40@andrew.cmu.edu

Lorrie Faith Cranor
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA, 15213
lorrie@cs.cmu.edu

ABSTRACT

Displaying website privacy policies to consumers in ways they understand is an important part of gaining consumers' trust and informed consent, yet most website privacy policies today are presented in confusing, legalistic natural language. Moreover, because website privacy policy presentations vary from website to website, policies are difficult to compare and it is difficult for consumers to determine which websites offer the best privacy protections. The Platform for Privacy Preferences (P3P) addresses part of the problem with natural language policies by providing a formal, machine-readable language for expressing privacy policies in a manner that is standardized across websites. To address remaining problems, an automated tool must be developed to read P3P policies and display them to users in a comprehensible way. To this end, we have developed a P3P policy presentation tool based on the Expandable Grid, a visualization technique for displaying policies in an interactive matrix. In prior work, the Expandable Grid has been shown to work well for displaying file permissions policies, so it appears to hold promise for presenting online privacy policies as well. To evaluate our Expandable Grid interface, we conducted two user studies, an online study with 520 participants and a laboratory study with 12 participants. The studies compared participants' comprehension of privacy policies presented with the Grid interface with their comprehension of the same policies presented in natural language. To our surprise, comprehension of policies was, for the most part, no better with the Grid interface than with natural language. We describe why the Grid interface did not perform well in our study and discuss implications for when and how the Expandable Grid concept can be usefully applied.

Categories and Subject Descriptors: H.5.2 [Information Interfaces and Presentation]: User Interfaces; K.4.1 [Computers and Society]: Public Policy Issues–Privacy

General Terms: Experimentation, Human Factors, Standardization

Keywords: privacy, P3P, Expandable Grid, policy, user interface

1. INTRODUCTION

Presenting website privacy policies to consumers in a clear and concise manner is important from at least two perspectives. First, from the perspective of consumer protection, fair information practice principles require that consumers be informed of how websites will use consumers' personal information. The United States Federal Trade Commission (FTC) names notice/awareness, i.e., giving consumers notice of an entity's information practices, as the first of its five core principles of privacy protection [9]. Second, from the perspective of websites trying to gain consumers' trust, providing clear notice of privacy practices may help allay consumer concerns about misuse of their personal data [2]. Moreover, for e-commerce websites, consumers may be willing to pay a premium if presented with a prominent display of consumer-friendly privacy practices [16].

However, despite the importance of clear presentation, privacy policies are usually presented in legalistic, convoluted language [13], and are often written at a college or higher reading level [1]. Researchers, industry groups, and privacy advocates have proposed other methods for presenting privacy policies, but there is no consensus on an effective presentation format. To make matters worse, presentations vary from website to website; thus, from the consumer's point of view, every privacy policy is as hard to read as the last one, and it is very difficult to compare privacy policies across competing websites.

The Platform for Privacy Preferences (P3P) was designed to address some of the drawbacks to natural language privacy policies [6, 7]. P3P is an eXtensible Markup Language (XML) based machine-readable language for expressing website privacy policies. It enables websites to specify policies in a uniform manner that can be read and presented by user agents, such as a Web browser or a policy-display ap-

WPES'08, October 27, 2008, Alexandria, Virginia, USA.
Copyright 2008 ACM 978-1-60558-289-4/08/10 ...$5.00.

plication like Privacy Bird [8]. Many websites provide P3P policies; a 2007 study found that 28% of the top 75 dot-com domains had been P3P enabled [5]. The P3P specification does not specify a presentation format.

We introduce the P3P Expandable Grid, an interactive format for presenting P3P policies to website visitors. The P3P Expandable Grid is based on the concept of the Expandable Grid, an information visualization technique for displaying policies [14]. Reeder et al. applied the Expandable Grid concept to a user interface for displaying and authoring file permissions policies. They showed that their Expandable-Grid-based interface was very effective for a wide range of file permissions policy-authoring tasks. Their result suggests that the Expandable Grid idea might work similarly well for displaying P3P policies. We designed and implemented the P3P Expandable Grid to see if the Expandable Grid would fulfill its promise as a P3P policy presentation format.

We evaluated the P3P Expandable Grid in two studies, one conducted on the Web and one conducted in our lab. In the Web-based study, 520 participants viewed website privacy policies either in natural language or in the P3P Expandable Grid and were asked comprehension questions about the policies. Participant performance was poor in both conditions, and, to our surprise, was generally worse with the P3P Expandable Grid. To explain this unexpected result, we conducted a lab study in which we collected detailed video and think-aloud audio data from 12 participants who viewed the same policies in both the natural language and P3P Expandable Grid formats and answered the same comprehension questions. We identified specific usability problems with the P3P Expandable Grid as a tool for displaying privacy policies to consumers. These usability problems, such as lack of a focal point, unintuitive organization of policy elements, and confusing icons and terminology, suggest both specific improvements that might make the P3P Expandable Grid an effective tool and general lessons as to when and how to apply the Expandable Grid concept to other policy domains. We discuss these suggested improvements and general lessons. We also note that while the current implementation of the P3P Expandable Grid did not work well for conveying privacy policies to consumers who had never seen the Grid before, it might be of use as an authoring tool for P3P experts or as the basis for a standardized privacy policy presentation.

2. RELATED WORK

Prior work has shown that natural language policies are difficult to read, with an average Flesch Grade Level of 14 (college level), even though only 27% of the US population has a college education [11]. A review of 60 financial privacy policies found they were all "difficult" to read or worse on the Flesch Reading Ease scale [10]. In addition to being difficult to understand, companies use textual ambiguity to obscure data collection practices. For example, they might state they perform a given practice "from time to time" to minimize its perceived importance [13].

The law firm Hunton & Williams has advocated layered online privacy policies [12]. Layered policies present a one-page summary with a link to more detailed information [3].

The Kleimann Communication Group studied Gramm-Leach-Bliley Act financial statements and developed improved designs for printed policies [15]. After extensive study of multiple presentations, they found a "...table design worked far better in helping consumers easily access, understand, and compare sharing practices" [15].

The Privacy Bird project uses P3P to create a standardized privacy report. The Privacy Bird report uses bulleted lists to summarize information and has a hide/expand feature so Internet users can focus on a high-level summary or drill down for full information about an online policy. Prior research has shown users like the Privacy Bird format and are able to answer comprehension questions about privacy policies using the format, but sometimes make errors because they fail to notice when information is collected on an opt-in or opt-out basis [8].

With the Expandable Grid, we used a table format as suggested by the Kleimann report. The Grid supports the ability to condense information by contracting the columns and rows, which is similar to the design goals of layered policies and Privacy Bird. Grid icons clearly show when information is opt-out or opt-in.

3. SYSTEM DESCRIPTION

Our system for displaying P3P policies, the P3P Expandable Grid, takes the elements of a P3P policy and maps them onto the conceptual framework of the Expandable Grid visualization. Here we describe the relevant elements of P3P, the Expandable Grid concept, and our implementation of the P3P Expandable Grid.

3.1 P3P

P3P defines a set of data practices in which an organization with a website could potentially engage. A specific organization's P3P policy states which of these practices the organization actually engages in. For example, one potential data practice is collecting a consumer's online contact information (e.g., email address) and sharing it with other companies to contact the consumer for marketing purposes. Another potential data practice is collecting a consumer's web navigation patterns and using them to improve a website's layout. Each potential data practice defined by P3P consists of three primary data-specific assertions: data category, recipient, and purpose.[1] P3P defines 17 data categories, 12 purposes, and six recipients. It also defines a hierarchy of 31 data elements, each of which can have numerous sub-elements (data elements are items within the data categories, such as email address or home phone number). Any combination of data element (or category), purpose, and recipient constitutes a potential data practice. A P3P policy consists of one or more statements, which are XML elements that each contain a group of data elements that are all to be handled similarly. A statement thus declares a set of the data practices an organization potentially engages in.

There are four distinct P3P policy outcomes that we name REQUIRED, NOT-USED, OPT-IN, and OPT-OUT. A P3P statement maps each of the potential data practices to a policy outcome. REQUIRED outcomes indicate data practices in which an organization engages; NOT-USED outcomes

[1] P3P defines six types of data-specific assertions, but to simplify this discussion, we leave out the infrequently-used non-identifiable assertion, the retention assertion (i.e., websites with P3P policies almost invariably retain data "indefinitely"), and the free-text human-readable consequence assertion.

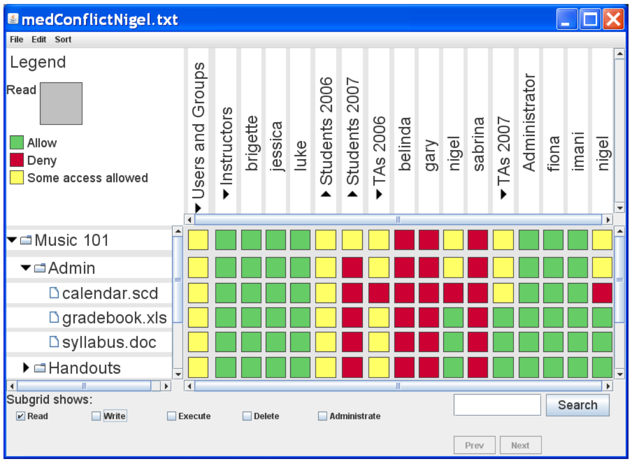

medConflictNigel.txt

File Edit Sort

Legend

Read

🟩 Allow
🟥 Deny
🟨 Some access allowed

	Users and Groups	Instructors	brigette	jessica	luke	Students 2006	Students 2007	TAs 2006	belinda	gary	nigel	sabrina	TAs 2007	Administrator	fiona	imani	nigel
Music 101	🟨	🟩	🟩	🟩	🟩	🟨	🟨	🟨	🟥	🟥	🟨	🟥	🟨	🟩	🟩	🟩	🟨
Admin	🟨	🟩	🟩	🟩	🟩	🟨	🟥	🟨	🟥	🟥	🟨	🟥	🟨	🟩	🟩	🟩	🟨
calendar.scd	🟨	🟩	🟩	🟩	🟩	🟨	🟥	🟥	🟥	🟥	🟥	🟥	🟨	🟩	🟩	🟩	🟥
gradebook.xls	🟨	🟩	🟩	🟩	🟩	🟨	🟥	🟨	🟥	🟥	🟩	🟥	🟩	🟩	🟩	🟩	🟩
syllabus.doc	🟨	🟩	🟩	🟩	🟩	🟨	🟥	🟨	🟥	🟥	🟩	🟥	🟩	🟩	🟩	🟩	🟩
Handouts	🟨	🟩	🟩	🟩	🟩	🟨	🟥	🟨	🟥	🟥	🟩	🟥	🟩	🟩	🟩	🟩	🟩

Subgrid shows:

☑ Read ☐ Write ☐ Execute ☐ Delete ☐ Administrate

Figure 1: Screenshot of the Expandable Grid interface for setting file permissions.

indicate data practices in which an organization does not engage; Opt-In outcomes indicate data practices in which an organization engages if a consumer requests to be subjected in the practice; and Opt-Out outcomes indicate data practices in which an organization engages unless the consumer requests not to be subjected to the practice.

In the preceding description, we have simplified some of the details of P3P. A full description of the language can be found in the book *Web Privacy with P3P* [7] or in the P3P specification [6].

3.2 The Expandable Grid

The Expandable Grid is a visualization technique for displaying a policy. Originally conceived for displaying file permissions policies in a file permissions authoring interface, the concept can be applied to any policy that maps elements from multiple dimensions to outcomes [14]. The Expandable Grid is a 2-dimensional matrix visualization with hierarchical trees of labels on each axis of the matrix and a grid in the middle showing a graphical representation of the policy outcomes corresponding to each pair of labels. For instance, the Expandable Grid interface for file permissions (see Figure 1) shows mappings from principals (users and groups) and resources (files and folders) to outcomes of Allow or Deny. Expandable trees representing hierarchies of users and files are displayed along the axes of the grid. The grid shows and hides cells as necessary in response to expansion or contraction of the trees – hence the name Expandable Grid. The tree along the vertical axis at the left of the interface shows the resources in a file system. The rotated tree along the horizontal axis at the top of the interface shows the principals. At the intersection of these two trees is a grid that shows the access each principal has to each resource. Grid cells each correspond to one principal and one resource. Green cells (which appear as a medium grey in greyscale) indicate access that is allowed, red cells (which appear dark grey in greyscale) indicate access that is denied, and yellow cells (which appear light grey in greyscale) indicate that items lower in one or both trees have a mixture of allowed and denied access.

3.3 P3P Expandable Grid design

P3P policies are more complex than typical file permissions policies. In a typical file permissions policy, there are three attributes—principal, resource, and action—but there are only a handful of possible actions, so actions are fairly easy to represent. A P3P policy contains three primary data-specific assertions, but all three are more complex than the action attribute in a file permissions policy. Moreover, P3P policies have four possible outcomes compared to two for file permissions policies, and P3P policies contain metadata outside the defined data practices. Because of the complexity of P3P policies, we knew that making a graphical representation of a P3P policy would be challenging.

Since P3P data practices are defined by the three primary data-specific assertions, we used hierarchical structures of the P3P data categories, purposes, and recipients as the labels along the axes of our P3P Expandable Grid. Our data hierarchy uses the data categories as high-level nodes, and places the P3P data element hierarchy under the categories as per the P3P specification [6]. We put an expandable tree representation of the data hierarchy along the vertical (left-hand) axis of our grid. P3P defines its purposes and recipients in flat lists of categories, but we added a layer of structure to simplify presentation. For example, we grouped the P3P "contact" and "telemarketing" purposes into one higher-level "marketing" category. Since, after putting data categories along the vertical axis, we had only one axis left to represent two hierarchies, we put both the recipient and purpose hierarchies next to each other on the horizontal (top) axis. The layout of the P3P Expandable Grid can be seen in Figures 2 and 3.

While putting recipient and purpose on the same axis would seem to restrict our ability to represent certain policies, namely those in which both recipient and purpose vary for the same data element, P3P itself has the same restriction, and requires multiple statements to represent such policies. Since our grid design can represent multiple statements, we are able to get around this restriction in the same way the underlying P3P language does. So, for example, a policy which allows email address to be collected by the company issuing the policy for the purpose of site administration, but also allows email address to be collected and shared with other companies for the purpose of marketing, will require multiple P3P statements.

The grid itself consists of squares at the intersections of each row representing a data element and each column representing a recipient or a purpose. The squares are colored according to the P3P policy outcome for each data practice defined by the combinations of data elements, recipients, and purposes. For each data element, the grid squares corresponding to recipients who use the data element or to purposes for which the element is used are colored teal; grid squares corresponding to recipients that do not use the data element or purposes for which the data element is not used are colored grey. We chose teal as a neutral color with no obvious conventional meaning; we were concerned that green or red might imply "good" or "bad" judgments about policy content. Grid squares at the intersection of non-leaf nodes of the hierarchies potentially represent multiple distinct outcomes of the leaf nodes beneath them, so they are colored with a teal-white gradient if there are both teal and grey squares beneath them, just as yellow squares in the file permissions Expandable Grid indicate a mixture of allow and

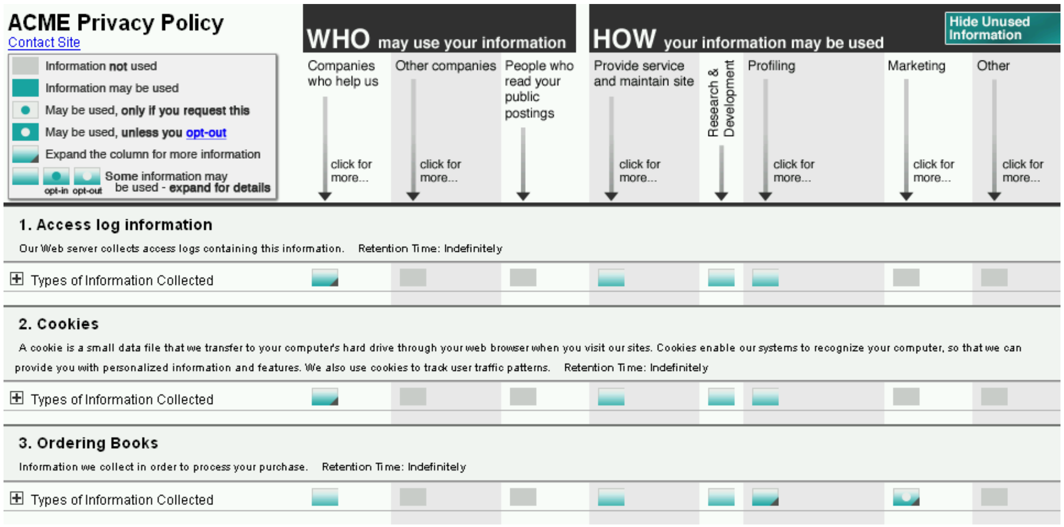
ACME Privacy Policy
Contact Site
Information not used
Information may be used
May be used, only if you request this
May be used, unless you opt-out
Expand the column for more information
Some information may be used - expand for details
opt-in opt-out
WHO may use your information
HOW your information may be used
Hide Unused Information
Companies who help us
Other companies
People who read your public postings
Provide service and maintain site
Research & Development
Profiling
Marketing
Other
click for more...

1. Access log information

Our Web server collects access logs containing this information. Retention Time: Indefinitely

⊞ Types of Information Collected

2. Cookies

A cookie is a small data file that we transfer to your computer's hard drive through your web browser when you visit our sites. Cookies enable our systems to recognize your computer, so that we can provide you with personalized information and features. We also use cookies to track user traffic patterns. Retention Time: Indefinitely

⊞ Types of Information Collected

3. Ordering Books

Information we collect in order to process your purchase. Retention Time: Indefinitely

⊞ Types of Information Collected

Figure 2: Screenshot of the P3P Expandable Grid in unexpanded form. The column headers and data hierarchies can be expanded and contracted to show more or fewer P3P elements.

deny access in the nodes beneath them. The different colored squares are shown in Figure 3.

Teal and grey squares cover the REQUIRED and NOT-USED outcomes of P3P policies, but not the OPT-IN and OPT-OUT outcomes, so we added a dot notation to indicate data practices for which the consumer was given the choice to opt into or out of. Squares are colored with a teal dot on a white background to indicate an OPT-IN outcome and a white dot on a teal background to indicate an OPT-OUT outcome. Dots over gradients indicate squares corresponding to non-leaf nodes under which there are some OPT-IN or OPT-OUT outcomes. Squares representing the OPT-OUT outcome can be seen in Figure 3.

All square designs may also contain a small black dog-ear symbol in the lower right-hand corner, indicating the cell may be clicked to expand its column. The dog-ear symbol can be seen on some of the squares in Figure 3.

Metadata associated with each P3P statement is contained in a header box of text above the data hierarchy. Multiple statements with header boxes and unexpanded data hierarchies can be seen in Figure 2.

Each data hierarchy is initially displayed in its fully collapsed form, i.e., only the root node of the data hierarchy is shown, and a user has to click to expand the root node in order to see data categories and elements (see Figure 2). The recipient and purpose hierarchies are also initially shown in their fully collapsed states.

Metadata associated with the entire P3P policy, such as information on opt-out mechanisms and organizational contact information, is in two places: a legend in the upper-left corner of the grid display contains an opt-out link (visible in Figures 2 and 3), and text below all of the data statements shows organizational contact information and any additional policy metadata (not shown in the figures).

4. METHODOLOGY

We conducted two user studies to compare the P3P Expandable Grid format for presenting privacy policies to the natural language format. Participants in our studies answered policy comprehension questions using either a privacy policy written in natural language or the same privacy policy written in P3P and presented with the P3P Expandable Grid.

4.1 Web-based user study

The first of our user studies was conducted over the Web. We used a between-participants design with two factors: presentation format and policy length. Format had two levels: Grid and natural language, and length had three levels: short, medium, and long, so there were a total of six conditions.[2] We assigned participants randomly to conditions.

4.1.1 Participants

We posted advertisements to a variety of online forums to recruit participants to complete our study. Participants received entry into a drawing for a $250 Amazon.com gift certificate. Online advertising forums included the Craigslist classified ad site, sweepstakes websites, mailing lists, and personal networks. We also purchased Google adwords, although fewer than ten participants came through Google ads. The variety of advertising venues we used was intended to attract our desired demographic of the general class of Web users. We recruited 786 participants to start the study, of which 520 completed the study and 266 dropped out.

4.1.2 Policies used

We presented the same privacy policy to participants in both the Grid and natural language conditions. The policy we chose was a real privacy policy from a major publisher. We chose this policy for the following reasons:

- It was real, and thus representative of a privacy policy a consumer might encounter in practice;
- It was published in both natural language and P3P versions;
- The P3P version had multiple data statements, and we wanted to be able to test the P3P Expandable Grid's ability to show a policy with multiple statements.

[2] This study was part of a larger study that included two other formats, layered natural language and Privacy Finder, in the comparison, but we exclude those formats from our discussion here; those results will be presented elsewhere.

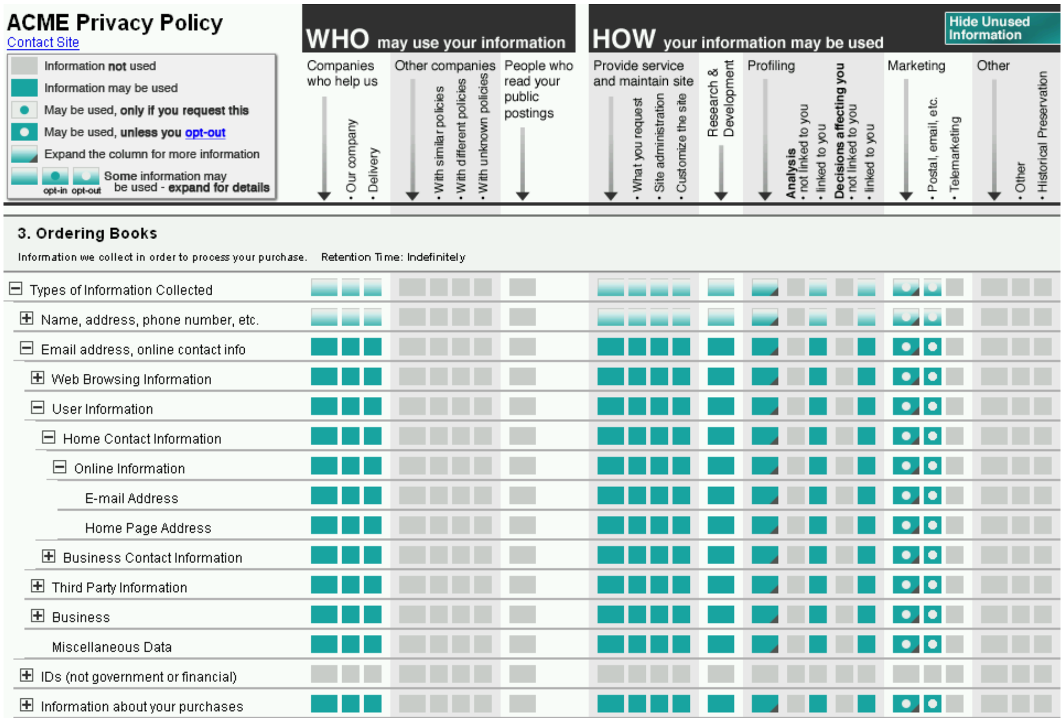

ACME Privacy Policy

Contact Site

⬜	Information **not** used
🟩	Information may be used
•	May be used, **only if you request this**
●	May be used, **unless you opt-out**
▤	Expand the column for more information
◩ • ● (opt-in, opt-out)	Some information may be used - **expand for details**

3. Ordering Books

Information we collect in order to process your purchase. Retention Time: Indefinitely

WHO may use your information
HOW your information may be used
Hide Unused Information

	Companies who help us			Other companies				People who read your public postings	Provide service and maintain site				Research & Development	Profiling					Marketing			Other		
	↓	• Our company	• Delivery	↓	• With similar policies	• With different policies	• With unknown policies	↓	↓	• What you request	• Site administration	• Customize the site	↓	↓	**Analysis** • not linked to you	• linked to you	**Decisions affecting you** • not linked to you	• linked to you	↓	• Postal, email, etc.	• Telemarketing	↓	• Other	• Historical Preservation
⊟ Types of Information Collected	▤	▤	▤	⬜	⬜	⬜	⬜	⬜	▤	▤	▤	▤	▤	◩	⬜	▤	⬜	▤	◩●	▤●	⬜	⬜	⬜	⬜
⊞ Name, address, phone number, etc.	▤	▤	▤	⬜	⬜	⬜	⬜	⬜	▤	▤	▤	▤	▤	◩	⬜	▤	⬜	▤	◩●	▤●	⬜	⬜	⬜	⬜
⊟ Email address, online contact info	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊞ Web Browsing Information	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊟ User Information	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊟ Home Contact Information	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊟ Online Information	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
E-mail Address	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
Home Page Address	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊞ Business Contact Information	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊞ Third Party Information	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊞ Business	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
Miscellaneous Data	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜
⊞ IDs (not government or financial)	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜
⊞ Information about your purchases	🟩	🟩	🟩	⬜	⬜	⬜	⬜	⬜	🟩	🟩	🟩	🟩	🟩	◩	⬜	🟩	⬜	🟩	◩●	🟩●	⬜	⬜	⬜	⬜

Figure 3: Screenshot of the P3P Expandable Grid in expanded form.

We changed references to the publisher's company name in the policy to "Acme" to avoid any associations with the real company from which the policy was taken. Since we were varying policy lengths, we wanted to present long, medium, and short versions of the policy. We used the whole real P3P policy, consisting of eight P3P data statements, as the long version, and we eliminated some of the statements from the whole policy to create a two-statement medium version and a one-statement short version. We also created medium and short versions of the natural language policy that corresponded to the medium and short versions of the P3P policy. Since the real natural language policy organized its content into paragraphs corresponding to the P3P policy's data statements, paring down the natural language policy to match the medium and short P3P policies was simply a matter of eliminating the paragraphs that corresponded to the eliminated data statements.

4.1.3 *Questions used*

We asked seven multiple choice comprehension questions about the privacy policies:

1. *Telemarketing:* Will Acme collect your home phone number and use it for telemarketing?
2. *Cookies:* Does the Acme website use cookies?
3. *Marketing Email:* Does this privacy policy allow Acme to put you on an email marketing list?
4. *Opt-out:* How can you remove yourself from Acme's email list?
5. *Share:* Does this privacy policy allow Acme to share your email address with a marketing company that might put you on their email marketing list?
6. *SSN:* Does the Acme website collect your Social Security number?
7. *Encryption:* If you send your credit card number to Acme do they keep it encrypted to prevent data theft?

For all questions except the Opt-out question, the multiple choice answers were "Yes," "No," or "The policy does not say." For the Opt-out question, multiple choice answers were "Call Acme's customer service," "Send email to Acme," "Click a link to opt out," "You cannot make any choices," and "The policy does not say," and participants were asked to check as many answers as applied.

The Telemarketing question was always presented first and served as a question to allow participants to gain some familiarity with the presentation format they were using. We excluded the Telemarketing question from data analysis.

We took the next four questions, Cookies, Marketing Email, Opt-out, and Share, from Cranor et al. [8], who found that they reflect common consumer privacy concerns.

We designed the last two questions, SSN and Encryption, specifically to test aspects of the Grid. The SSN question tested whether participants could determine that the company does not engage in a particular practice. In the policy we selected, Social Security number is not collected. One advantage of the Grid over natural language is that it specifically shows data practices in which an organization does not engage through the use of grey boxes. Natural language policies generally do not discuss practices in which the organization does not engage, so it is often difficult to determine decisively from a natural language policy that an organization does not engage in a practice. The Encryption question asked for information that actually was not covered in the privacy policy; we expected participants using both formats to struggle to answer the question, but expected it might be easier to see in the Grid layout that encryption is a security feature, not a data practice defined by P3P.

4.1.4 *Subjective satisfaction questions*

After participants completed comprehension questions, we asked participants to rate their agreement with a series of statements regarding their subjective satisfaction with the privacy policy presentation. Example statements include "Finding information in Acme's privacy policy was a pleasurable experience" and "I feel confident in my understanding of what I read of Acme's privacy policy."

4.1.5 Procedure

When participants visited our study website, they were randomly assigned to one of the two formats and one of the three policy lengths. The two formats and three policy lengths made for six total conditions, so we implemented random assignment by assigning each subsequent participant to the next condition in a repeating sequence of the six conditions. After being assigned to a condition, participants were presented a Web page with a comprehension question and a multiple choice form for completing the question in a frame at the top of the page, and a privacy policy presentation in the assigned format and length in a frame beneath the comprehension question. All participants were presented the Telemarketing question first. After participants answered the Telemarketing question, they were presented the remaining six questions in random order to guard against learning and sequencing effects. Following completion of the comprehension questions, participants were asked the subjective satisfaction questions.

4.1.6 Data collected

For each participant, we recorded their multiple choice answers to comprehension questions, the time it took them to answer each question, whether or not they finished the comprehension questions, when they expanded rows or columns in the Grid, and their answers to the subjective satisfaction questions. We scored each multiple choice answer as correct or incorrect to compute accuracy rates for each question and each condition.

4.2 Lab-based user study

We followed up the Web-based study with a similarly designed study in our laboratory that allowed us to collect detailed qualitative data to help explain the quantitative results of the Web-based study. We recruited 12 participants via a local university newsgroup. Participants answered the same seven comprehension questions used in the Web-based study, but policy presentation conditions were somewhat different. Presentation format was a within-participants variable, so each participant answered all seven comprehension questions twice, once using the Grid format and once using the natural language format. Participants did not view the same policy in both formats, however; if they viewed the short policy in the Grid, they viewed the long policy in natural language, and vice versa. We dropped the medium policy length and used only the short and long policy lengths. We alternated the order in which the two formats were presented, so that six of the participants viewed a policy in the Grid first, and six viewed a policy in natural language first. We asked participants to think aloud as they worked, and we recorded screen video and think-aloud audio for all participants.

5. RESULTS

For the Web-based study, we measured accuracy rate and mean time to question completion for each of the six comprehension questions and for each of the six conditions. Accuracy rate is the proportion of participants who correctly answered a comprehension question. Before computing accuracy rate and time to question completion, we checked the data for participants who may have gamed the study by clicking as quickly as possible through the comprehension questions without putting effort into answering correctly. We eliminated from analysis those participants who finished faster than two standard deviations below the mean of the log-transformed time to question completion data for three or more questions. We eliminated seven participants for this reason, including six Grid participants and one natural language participant.

In addition to accuracy rate and mean time to question completion, we computed scores on the subjective satisfaction questions.

For the lab-based study, we analyzed video and audio for evidence of usability problems with the Grid.

5.1 Accuracy results

We used logistic regression to test for an overall difference in accuracy rates across all six comprehension questions between the Grid (overall accuracy = 0.59) and the natural language (overall accuracy = 0.68) formats. We used format as the single factor in our logistic regression model. The model gave an intercept of 0.27 and a coefficient for format of 0.25. A Wald test of the hypothesis that the format coefficient was not equal to 0 was significant at the 0.05 level ($Z = 3.62, p < 0.001$), suggesting that there is an effect of format on accuracy. The sign of the coefficient, as well as the actual overall accuracy rates, suggests that over all questions, participants were more likely to correctly answer questions using the natural language format.

To follow up the logistic regression and gain a more detailed picture of how the Grid compared to the natural language format, we performed 18 pairwise comparisons of the Grid accuracy rate with the natural language accuracy rate for each question and for each policy length. We used one-sided Fisher's exact tests of the hypothesis that the Grid accuracy rates were higher than natural language accuracy rates. Since our experimental goal was to determine whether the Grid format was superior to the natural language format, we did not test the hypothesis that the natural language accuracy rates were higher than the Grid accuracy rates, even though, as it turns out, the natural language accuracy rates were in most cases higher. We corrected for multiple testing using the Benjamini-Hochberg method, which gave an adjusted per-test $alpha$ of 0.00008.

These pairwise comparisons showed that the Grid gave significantly higher accuracy rates for only one question, the SSN question, at the short and medium policy lengths. The accuracy rate data is shown in Table 1.

5.2 Time to question completion results

To test the hypothesis that the Grid format would lead to faster performance on comprehension questions, we computed mean time to question completion for each question at each policy length and for each format. Mean time to question completion for a comprehension question included completion times only for those participants who correctly answered the question, since we are interested in the time it takes to correctly comprehend a policy. We also excluded data points that were more than three standard deviations above or below the mean of the log-transformed time to

Table 1: Summary of accuracy rate (a) and mean (with standard deviation) time to question completion (t) data in seconds for the P3P Expandable Grid (a_G, t_G) and for natural language (a_{NL}, t_{NL}), for each comprehension question at each policy length. Statistically significant results are shown in bold.

Task	Short				Medium				Long			
	a_G	a_{NL}	$t_G(sd)$	$t_{NL}(sd)$	a_G	a_{NL}	$t_G(sd)$	$t_{NL}(sd)$	a_G	a_{NL}	$t_G(sd)$	$t_{NL}(sd)$
Cookies	0.90	0.99	**7(8)**	**9(6)**	0.92	0.98	15(14)	9(7)	0.96	0.99	14(14)	11(9)
Marketing Email	0.84	0.78	21(25)	14(16)	0.76	0.69	35(39)	15(14)	0.70	0.65	31(61)	30(46)
Opt-out	0.12	0.66	81(79)	33(18)	0.15	0.62	72(56)	36(20)	0.15	0.67	72(70)	35(18)
Share Email	0.21	0.47	53(43)	29(25)	0.17	0.54	33(42)	26(22)	0.44	0.54	47(103)	35(45)
SSN	**0.56**	**0.27**	28(28)	17(16)	**0.71**	**0.27**	39(43)	11(8)	0.66	0.49	47(48)	18(15)
Encryption	0.85	0.91	66(56)	22(15)	0.77	0.91	102(127)	27(20)	0.86	0.87	120(232)	44(55)

question completion data for any one question, policy length, and format. We assumed that such data points were the result of participants taking a break from the study in the middle of a question or rushing through the question, rather than merely the result of exceptionally slow or fast performance on the question. We excluded 11 data points (of 3078, which is 520 participants minus 7 eliminated for gaming the study times six questions), distributed across all six conditions.

An ANOVA comparing the overall hypothesis that the mean time to question completion for all questions was different between the Grid ($M = 44.7, sd = 90.0$) and natural language ($M = 23.2, sd = 27.0$) formats was significant at the $\alpha = 0.05$ level ($F(1, 1966) = 31.055, p < 0.001$). The direction of the difference in means suggests the natural language format led to faster performance at completing comprehension questions, and the results of the ANOVA show this result to be statistically significant.

We followed up the ANOVA with 18 pairwise comparisons of the mean time to question completion of each format for each question and policy length. We used one-sided t-tests to test the hypothesis that Grid mean time to question completion was less than natural language mean time to question completion. As with accuracy, we were not concerned with testing for significantly better performance by the natural language format since our experimental goal was to determine when the Grid format is an improvement on the natural language format. We corrected for multiple testing using the Benjamini-Hochberg method, which gave an adjusted per-test α of 0.004. The time to question completion data is shown in Table 1.

5.3 Subjective satisfaction results

We compared participants' subjective satisfaction with each presentation format. Participants' responses on the subjective satisfaction questions were on a scale from 1 (the strongest negative attitude) to 7 (the strongest positive attitude). We computed the mean response across all questions for participants who used the Grid ($M = 2.8, sd = 1.6$) and natural language ($mean = 3.8, sd = 1.5$). Two t-tests showed both the Grid ($t = -32.9(1899), p < 0.0001$) and natural language ($t = -5.8(2180), p < 0.0001$) means to be statistically significantly less than 4.0, the neutral score, at the $\alpha = 0.05$ level. This result suggests that participants in both formats had negative attitudes toward reading privacy policies. Another t-test showed the difference in means between the Grid and natural language formats to be significant at the $\alpha = 0.05$ level ($t = -20.9(3927), p < 0.0001$), suggesting attitudes were more positive toward the natural language format.

5.4 Lab-based study results

The objective of the data analysis for the lab-based study was to determine reasons why participants answered comprehension questions incorrectly using the Grid. To this end, we performed three analyses of the videos and think-aloud audio from the study. The three analyses together revealed eight usability problems that led to participants incorrectly answering comprehension questions. These analyses and usability problems are described in this section.

5.4.1 *Analysis of where answers were found*

In instances where participants incorrectly answered questions, we attempted to determine where in the Grid presentation the participant found the incorrect answer. It was sometimes possible based on what was on the screen, what participants had recently clicked on or scrolled to, and what participants were saying to pinpoint specific Grid squares, rows, or metadata from which they determined answers. Because this process was somewhat subjective, we used two raters and only included data from cases in which both raters agreed on where the participant had found an answer.

This analysis revealed four causes of incorrect answers. The causes were:

- *Multiple statements.* Answering comprehension questions for medium and long P3P policies was difficult because the policies contained multiple P3P statements. For the Cookies and Acme Email questions, the answers were only contained in a subset of the multiple statements, and for the Share Email and SSN questions, it was necessary to check all the statements. We observed five instances of participants finding an answer in the wrong statement or failing to check all statements.
- *Metadata.* Policy metadata that was out of the graphical region of the Grid was hard to find. Not surprisingly, participants seemed to focus on the graphical portion of the Grid presentation. In the Opt-out question, which required that participants read policy metadata, participants sometimes missed all or part of the necessary metadata to answer the question. We observed seven instances of such problems.
- *Terminology.* Participants generally seemed confused by P3P concepts and terminology. For example, we observed two instances of participants looking in the "Companies who help us" recipient column for the Share Email task, when they should have been looking in the "Other companies" recipient column. In the P3P policy, "Companies who help us" indicates companies who receive customer data to help complete orders but

not for marketing purposes. "Other companies" indicates companies that might receive customer data for marketing purposes. Other examples P3P terms we found, informally, to be confusing included the "Opt-in" and "Opt-out" outcomes and the term "Profiling" to cover P3P's `pseudo-analysis`, `pseudo-decision`, `individual-analysis`, and `individual-decision` purposes, all of which are related to creating profiles of users.

- *Two dimensions, one axis.* Because P3P policies have three primary data-specific assertions, it really requires a three-dimensional table to represent a whole policy. In our design, we juxtaposed the recipient and purpose assertions together on one axis. This juxtaposition caused errors. In the Share Email question, participants had to find a row corresponding to "Email address" and then notice that the "Other companies" recipient column had a grey square. However, we observed two instances of participants thinking they found the answer in the "Marketing" purpose column. Had the purposes been separate from recipients, this error likely would not have happened.

5.4.2 Analysis of confusion statements.

We analyzed the think-aloud audio for instances in which participants expressed confusion. Two raters listened to the think-aloud audio and noted statements that, in their judgment, indicated confusion. Examples of such statements included:

- "The policy itself is not that – very clear at all."
- "Yeah, I'm just confused right now what I'm looking at. OK, so there's these color-coded things..."
- "I'm not sure how to find this, I'm just kind of looking around."

We took the union of all confusion statements identified by the two raters, identified three candidate problem categories that might be indicated by the statements, and then had two raters place the statements into the three categories. The two raters identified 30 total confusion statements. The statements were distributed across 11 of the 12 participants and five of the six tasks, so it does not appear that any one task or any one participant was responsible for a disproportionate number of the statements. When categorizing the statements, the raters agreed on categories for 21 of the 30 statements and agreed that six of the statements fit into none of the three categories. The raters disagreed on the categorization of three of the statements. Statements on which raters disagreed or could not place into a category were eliminated from analysis.

Problem categories include:

- *No focal point.* Perhaps the most common problem with the Grid was its lack of a clear place to start looking at it. The Grid contains a great deal of information, and can look visually cluttered. There is so much to see that participants often missed the information they needed to answer questions correctly.
- *Multiple statements.* As we found in the previous analysis, participants were confused by multiple P3P statements.
- *Confusing icons.* The Grid has fifteen distinct symbols that can be displayed in Grid squares, and participants sometimes did not understand what they all meant.

5.4.3 Analysis of Grid expansions.

Analyzing data from both the Web-based and lab-based studies on when participants expanded the Grid, we found two problems: first, that some participants seemed not to realize the Grid was interactive, and second, that many participants did not understand the P3P data hierarchy.

Non-expansion.

We observed participants' expansions of the Grid's rows and columns. From the expansion data we recorded in the Web-based study, we observed that 35 of 241, or 14.5%, of Grid participants never expanded the Grid. In the lab-based study, we observed that one of 12 Grid participants never expanded the Grid. Some participants seem not to have realized the Grid is interactive.

Excess expansion.

We informally observed that some participants were not familiar with the P3P data hierarchy and had to do quite a bit of searching to find some of the items they were looking for. For example, participants looking for "Social Security number" in the SSN task did not know that P3P has a "Social Security numbers and government IDs" category. They would scan the list of categories for one that seemed relevant to Social Security numbers, find "Social and economic categories," and click on that. Only later would they find what they were looking for under "Social Security numbers and government IDs," at the bottom of the list of data categories.

To quantitatively confirm that lack of understanding of the P3P hierarchy was problem, we counted the number of excess expansions of the P3P data hierarchy in the lab-based study. We defined an "excess expansion" the expansion of a Grid row or column that only revealed Grid squares that were irrelevant to answering the question at hand. Thus, if a participant expanded "Social and economic categories" during the SSN task, this expansion was an excess expansion. We observed 13 excess expansions distributed across five of the six questions and seven of the 12 participants.

6. DISCUSSION

Our results strongly suggest that the P3P Expandable Grid, as currently implemented, is not an effective means for presenting privacy policies to Web users. Except in two cases (the SSN question and the Cookies question), participants using the Grid performed no better and no faster in correctly answering comprehension questions than participants using natural language. Moreover, subjective satisfaction scores show participants strongly disliked the Grid.

Policy length did not have the effect we expected on our results. We expected the Grid to be a more scalable format than natural language, but, if anything, we witnessed the opposite effect. Multiple statements in the policy represented in the Grid led to the problem of checking the wrong statement or failing to check all relevant statements.

Our results show two bright spots for the Grid. First, participants using the Grid did perform better, as expected, on the SSN task, in which they were asked to determine that a website did *not* engage in a certain data practice. In real tasks, users may want to reassure themselves that a company does not engage in some objectionable practice. Second, participants answered the Cookies question faster for the short policy using the Grid. While we cannot conclusively state the reason for the faster performance on the Cookies task, informal observations from the lab-based study suggest some participants visually searched for the word "Cookies," and found it faster in the Grid because there were fewer words to search through.

However, our results are mostly negative with respect to the Grid, and this is somewhat surprising given how well the Expandable Grid concept has been shown to work in a user interface for another policy domain, file permissions [14]. Our lab-based study suggests eight reasons why the Grid was difficult to use:

1. *No focal point.* The Grid is visually busy, so it is hard to know where to start a visual search for information.
2. *Difficulty with hierarchy.* Users are not familiar with the P3P data hierarchy and do not know how to find relevant items in it.
3. *Multiple statements.* Policies with multiple P3P statements may lead users to find information in the wrong statement or to fail to check all relevant statements.
4. *Metadata.* Policy metadata is difficult to find.
5. *Confusing icons.* The meaning of some of the 15 possible icons is not apparent to some users.
6. *Terminology.* P3P privacy policy terminology is confusing and unfamiliar to many users.
7. *Two dimensions, one axis.* Juxtaposing two dimensions on one axis was confusing to users.
8. *Non-expansion.* Some users never realized the Grid is interactive and can be clicked to expand it.

6.1 Lessons for applying the Expandable Grid concept

The usability problems we observed suggest several lessons for how to design presentations of policies based on the Expandable Grid concept. We list these lessons in this section.

Provide a starting point such as a search bar.

The most common problem and probably the biggest impediment to finding answers in the Grid was the vast amount of visual information with no clear starting point. User interfaces using the Expandable Grid concept should provide a starting point; a search feature may serve as a starting point. A search bar was part of the successful file permissions Expandable Grid design [14].

Provide a summary display.

The problems with lack of a focal point, difficulty with the P3P data hierarchy, and multiple statements could all be mitigated by simplifying the display of P3P policies. To this end, a simple summary display in which all statements are merged into one would be helpful. Also helpful would be a summary showing only data practices of the highest concern to consumers, such as sharing of health data or sharing contact information data for marketing purposes.

Use short labels.

The usability problems with terminology and metadata suggest some policies may not be well suited to display in an Expandable Grid. The problem of finding terminology to describe privacy concepts to Web users is known to be difficult. The P3P1.1 specification acknowledges the problem of finding concise descriptions for privacy concepts and offers some such labels for P3P elements [4]. In our case, the problem was exacerbated by trying to find short terms that could be used as labels in the Grid. The Expandable Grid may be best suited to showing policies in which the concepts can be summarized with short labels on the axes (e.g., in file permissions, users have names that are generally around eight characters long and files have names that are usually less than 30 characters). We tried using short labels and provided longer descriptions when users moved their mouse over the labels, but participants in our lab study still seemed confused by the concepts. There might be better labels for P3P concepts than those we used, but natural language may be the most effective means for explaining nuanced policy concepts to users who have never encountered those concepts before. Natural language may also be more effective for presenting policies with large amounts of metadata that cannot be put into the matrix format required by the Grid.

Answer common questions in one place.

Multiple statements in P3P policies make it necessary sometimes to check multiple grid squares to find the answer to a simple, common question. Expandable Grid interfaces should be designed so that common questions can be answered by looking in one place. A summary view in which the contents of multiple statements are merged, as described above, could provide one place to look for answers to common questions.

Place one dimension per axis.

A tabular representation of a P3P policy requires three dimensions, one for each of the three P3P data-specific assertions. To fit three dimensions into a two-dimensional display, we tried juxtaposing two dimensions together on the same axis. The approach did not work. Some solution is needed to fit policies of more than two dimensions into a Grid presentation; while we do not know what solution is right, juxtaposing two dimensions together seems to be wrong. A better solution might be to require participants to choose a fixed cross-section of the policy to look at, e.g., to select "Marketing" as the purpose and then view a two-dimensional grid of data-by-recipient.

Provide plenty of explanation for icons.

It is hard to help people understand a large number of icons. A legend helps, but may not be enough, since it draws the eyes away from the icons of interest. Help might be provided by showing the meaning of an icon when the mouse is moved over it. Also, reducing the number of icons needed to display a policy would help.

Emphasize or eliminate interactivity.

Users may have difficulty learning that a Grid presentation is interactive. We provided many cues of interactivity, including labels that read "click for more," mouse-over highlights on rows, and "+" symbols to indicate expandability. Still, many study participants did not realize the Grid is interactive. There might be even more indications we could provide to emphasize the Grid's interactivity, but another solution might be to eliminate interactivity altogether with a more compact policy representation that could fit on the screen without requiring expansion.

6.2 Limitations of our studies

We tested the P3P Expandable Grid on participants who had never seen the Grid before and had probably never encountered P3P before. While the Grid did not work well for displaying privacy policies to this class of users, it may be well suited to other use cases we did not test.

In particular, the Grid may work well as a tool for P3P experts authoring a P3P privacy policy. Indeed, our experience has been that we, as P3P experts, find it easier to understand the contents of a P3P policy by viewing it in the Grid's tabular form than in other common formats, such as a list of rules or raw XML. The Grid could be made interactive, so that authoring a P3P policy would be a simple matter of clicking on Grid squares to edit the policy. P3P policy authoring is currently done in raw XML or using somewhat unwieldy tools like the IBM P3P Policy Editor [7].

Another use case to which the Grid may be better suited is as a standardized privacy policy presentation. Participants in our study had never seen the Grid format before. If Web users saw privacy policies presented in a consistent format, such as the Grid, repeatedly across all websites, they might become more familiar with that format over time and ultimately find it easy to read. A standardized format would also help users compare privacy policies between websites. In contrast, comparing natural language policies across websites can be extraordinarily difficult.

We tested the P3P Expandable Grid presentation format against the natural language presentation format for a single organization. We believe the organization we chose may have an exceptionally well-written natural language privacy policy. The Grid may perform better against average or poorly written privacy policies. In future work, we will test with additional organizations' policies.

7. CONCLUSION

We developed a means for graphically presenting website P3P privacy policies based on the Expandable Grid concept and found that it did not generally improve Web users' comprehension of privacy policies compared to a natural language policy presentation. However, the Grid did perform well at indicating practices *not* allowed by a policy, and it may hold promise for allowing P3P experts to navigate and author policies. A simplified Grid presentation of P3P policies may be an effective means for presenting policies even to general Web users in a standardized format. In future work, we will continue efforts toward designing such a simplified P3P policy presentation.

8. ACKNOWLEDGMENTS

The authors would like to acknowledge Keisha How for technical work on the Web-based user study and Tony Poor and Danielle Chang for help analyzing data from the lab-based study. This work was supported in part by National Science Foundation Cyber Trust Grants CNS-0433540 and CNS-0627513, U.S. Army Research Office contract DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab, and the IBM OCR project on Privacy and Security Policy Management.

9. REFERENCES

[1] A. I. Anton, J. B. Earp, Q. He, W. Stufflebeam, D. Bolchini, and C. Jensen. Financial privacy policies and the need for standardization. *IEEE Security and Privacy*, 2(2):36–45, March/April 2004.

[2] M. Brown and R. Muchira. Taxonomy of conflicts in network security policies. *IEEE Communications*, 5(1):62–70, March 2004.

[3] Center for Information Policy Leadership. Ten steps to develop a multilayered privacy policy, 2007. Available at `http://www.hunton.com/files/tbl_s47Details%5CFileUpload265%5C1405%5CTen_Steps_whitepaper.pdf`. Accessed May 19, 2008.

[4] L. Cranor, B. Dobbs, S. Egelman, G. Hogben, J. Humphrey, M. Schunter, D. A. Stampley, and R. Wenning. The Platform for Privacy Preferences 1.1 (P3P1.1) specification. W3C Recommendation, November 2006. Available at `http://www.w3.org/TR/P3P11/`. Accessed May 19, 2008.

[5] L. Cranor, S. Egelman, S. Sheng, A. M. McDonald, and A. Chowdhury. P3P deployment on websites. *Electronic Commerce Research and Applications*, 50, 2008. To appear. Available at `http://lorrie.cranor.org/pubs/p3p-deployment.pdf`. Accessed February 26, 2008.

[6] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) specification. W3C Recommendation, April 2002. Available at `http://www.w3.org/TR/P3P/`. Accessed May 19, 2008.

[7] L. F. Cranor. *Web Privacy with P3P*. O'Reilly, Sebastopol, CA, 2002.

[8] L. F. Cranor, P. Guduru, and M. Arjula. User interfaces for privacy agents. *ACM Transactions on Computer-Human Interaction*, 13(2):135–178, June 2006.

[9] Federal Trade Commission. Privacy online: A report to congress, June 1998. Available at `http://www.ftc.gov/reports/privacy3/priv-23a.pdf`. Accessed February 26, 2008.

[10] M. Hochhauser. Lost in the fine print: Readability of financial privacy notices, July 2001. Available at `http://www.privacyrights.org/ar/GLB-Reading.htm`. Accessed May 19, 2008.

[11] C. Jensen and C. Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 471–478, New York, NY, USA, 2004. ACM.

[12] R. Lemos. Msn sites get easy-to-read privacy label. *CNET News.com*, 2005. Available at `http://news.com.com/2100-1038_3-5611894.html`. Accessed on May 19, 2008.

[13] I. Pollach. What's wrong with online privacy policies? *Communications of the ACM*, 50(9):103–108, September 2007.

[14] R. W. Reeder, L. Bauer, L. F. Cranor, M. K. Reiter, K. Bacon, K. How, and H. Strong. Expandable grids for visualizing and authoring computer security policies. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems(CHI 2008)*, New York, NY, 2008. ACM Press.

[15] Report by Kleimann Communication Group for the FTC. Evolution of a prototype financial privacy notice, 2006. Available at `http://www.ftc.gov/privacy/privacyinitiatives/ftcfinalreport060228.pdf`. Accessed May 19, 2008.

[16] J. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. In *The 6th Workshop on the Economics of Information Security (WEIS)*, 2008. Available at `http://weis2007.econinfosec.org/papers/57.pdf`. Accessed February 26, 2008.