Textured Agreements: Re-envisioning Electronic Consent

Matthew Kay, Michael Terry
HCI Lab, David R. Cheriton School of Computer Science
University of Waterloo
{mjskay, mterry}@cs.uwaterloo.ca

ABSTRACT
Research indicates that less than 2% of the population reads license agreements during software installation [7]. To address this problem, we developed *textured agreements*, visually redesigned agreements that employ *information layering*, *vignettes*, *sensationalism*, and *visual variety* to accentuate information and highlight its personal relevance. Notably, textured agreements accomplish these goals without requiring modification of the underlying text. A between-subjects experimental study with 84 subjects indicates these agreements can significantly increase reading times. In our study, subjects spent approximately 30 seconds longer on consent screens than the in control condition, where subjects spent only 7 seconds, on average. Furthermore, the study results indicate that the effects observed are not due to the novelty of the textured agreements' visual appearance alone, but rather, particular features of the designs. These results provide convincing evidence of the potential for textured agreements to positively impact software consent processes.

AUTHOR KEYWORDS
Informed consent, end-user license agreement, EULA

ACM CLASSIFICATION KEYWORDS
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
It is estimated that less than 2% of the population fully reads end-user license agreements (EULAs) [7]. However, these agreements contain important legal information and, increasingly, privacy policies. Like previous work (e.g., [5,6]), we consider the agreement process to be one of gaining *informed consent* and the agreements themselves to be a form of *consent agreement*.

Numerous factors contribute to the low reading rate of consent agreements. From the perspective of the user, the process of gaining consent does not directly contribute to the primary goal of *using* the software, and thus has low utility [5,6]. The content and presentation of material in consent agreements are also strong deterrents to reading them:

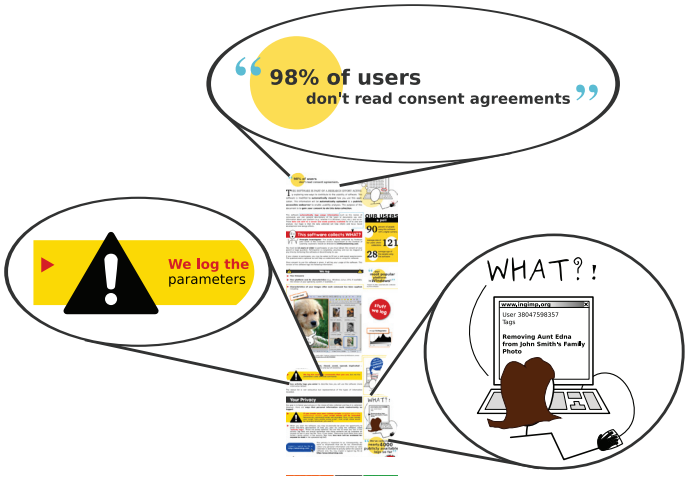

Figure 1. Textured agreements introduce information layering, vignettes with mini-narratives, sensationalism, and visual variety to compel users to read consent agreements.

The language used is often complex and filled with legalese [7,10], while the lengths of the agreements can be daunting [7].

Recognizing these issues, recent work in the HCI community has begun to investigate ways to improve the consent process in software systems. For example, Good *et al.* developed single-screen summaries to complement EULAs [7,8]. In a laboratory study, they found that the summaries resulted in significantly more users canceling the installation of spyware compared to a control condition [7]. The effectiveness of summaries in deterring the installation of spyware is promising, but it does not address the need to compel people to read the *full agreement*.

This paper introduces *textured agreements*, visually redesigned consent agreements that employ *information layering*, *vignettes*, *sensationalism*, and *visual variety* to accentuate information and highlight its *personal relevance*. For example, textured agreements use warning symbols to highlight terms of an agreement that may affect the user's privacy (Figure 1). Importantly, these techniques require no modification to the actual text of the agreement and can be adapted to a wide range of consent agreements.

To test the efficacy of textured agreements, we conducted a between-subjects experiment comparing five conditions: three textured agreements using minimal, moderate, and heavy application of the techniques; a summary condition to partially replicate Good *et al.*'s study [7]; and a control condition using only plain-text. The results of this study indicate that textured agreements are an effective means of compelling users to engage in the consent process. In particular,

Please cite as *Technical Report CS-2009-19, David R. Cheriton School of Computer Science, University of Waterloo, 2009*

moderate and heavy application of the techniques significantly increased reading time by approximately 30 seconds. Users also self-reported reading these agreements more than the control condition. Notably, the reading times were not significantly increased with merely decorative adornments added to the agreement, suggesting the effects observed in the moderate and heavy conditions are not due to the visual novelty alone.

The findings from this study are significant for several reasons. First, the positive effect of visual design techniques suggests they are a promising method for addressing the problem of low reading rates in software consent processes. Second, these techniques achieve their effectiveness by operating on the primary object of interest (the agreement), rather than creating secondary objects, such as summaries. Finally, textured agreements demonstrate improvements using only a small subset of visual design techniques, making it a promising technique by virtue of the wide range of additional visual design techniques that can complement those chosen for this work.

The study results also serve to confirm the effectiveness of summaries, but also caution against their use. More specifically, our results indicate that users read summaries, but at the cost of reading the actual consent agreement. This finding suggests the need to further develop the summary strategy so that its benefits are retained (e.g., preventing the installation of spyware), without it negatively impacting one's tendency to read the full consent agreement.

The rest of this paper describes related work in software-based consent agreements, our design methodology for improving the consent process, textured agreements, our experimental method for testing the agreements, and results and implications from the experimental study. We also provide perspectives from our local internal review board on the feasibility of deploying textured agreements in real-world situations.

BACKGROUND

Informed Consent: Goals and Challenges

As mentioned in the Introduction, we consider any software-based consent process to be an instance of gaining *informed consent* and the agreement itself to be a form of *consent agreement*. Software systems seek informed consent in a variety of contexts, such as when presenting a EULA. However, gaining *informed* consent is challenging because of the *low perceived value* of the process, *desensitization* and *habituation* to the activity, and the *ineffective presentation* of information. We describe each issue in turn.

In current systems, the consent process interrupts the user's primary task, requiring users to provide consent before the task can proceed [5,6,7]. Since the process blocks task progression, the optimal path is to provide consent as fast as possible so one can return to the primary task. Investing time is further discouraged by the simple fact that the consent process is one-sided: the end-user has no real power to influence the process [5,6]. These factors contribute to the process having *low perceived value* to users.

The frequency with which users encounter these consent agreements also adversely affects the consent process. In particular, repeated exposure can lead users to become *desensitized* and *habituated* to the agreements, which results in quick dismissal of the consent agreements when encountered [8].

Finally, many consent agreements overwhelm the user with complexity and length, making them difficult to navigate and fully comprehend [10]. For example, assessments of existing consent agreements indicate that they often require an advanced education to fully comprehend [10]. In essence the information is *ineffectively presented* to users.

Past research clearly indicates the failure of current systems to effectively address these challenges. It is estimated that less than 2% of the population regularly reads EULAs when installing software [7], and there is reason to believe few individuals read websites' privacy policies as well [10]. Research in usable security indicates that user apathy is also common when making assessments regarding the security risks posed by certain actions, such as installing third-party software or plug-ins [3]. However, a number of strategies have been developed in both industry and the research community to begin to address these issues.

Current Strategies for Improving the Consent Process

Current research aimed at improving the consent process has largely focused on addressing the problems of *habituation* and *ineffective presentation* of information.

Numerous systems have sought to prevent habituation by purposefully interrupting the normal flow of the consent process. For example, some systems rearrange the order of buttons on the interface so that users cannot learn a fixed path through the consent interface [3]. This technique has been shown to increase the likelihood that users take time to understand the information presented, making it a promising approach.

Noting the length and density of typical EULAs, Good *et al.* focused on improving the effectiveness of presenting information by creating single-screen *summaries* of EULAs [7,8]. In a study investigating their effect on the installation process, the researchers found that providing summaries before or after a EULA significantly reduced the number of installations of spyware [7].

These approaches are important steps in improving the consent process. However, we note that none of these techniques have explored ways to increase the *perceived value* of engaging in the consent process. We also note that the software consent process is largely conducted using the medium of *plain text*, with minimal use of typography or other visual design techniques. However, there is a rich repertoire of visual design techniques whose sole purpose is to suggest the value in taking notice of information. Countless visual techniques have also been developed to make communication fast and efficient. While a small number of these techniques have been explored for paper-based consent processes, such as financial privacy agreements [1,11] and medical informed consent [4,9], these strategies have not been examined in the

context of software consent. We turn now to our research that explicitly explores this approach.

RE-ENVISIONING THE CONSENT PROCESS

From the challenges identified in the previous section – low perceived value, habituation and desensitization, and ineffective presentation – we can develop a set of goals for software-based informed consent processes. In particular, an effective consent process will:

- *Compel* users to engage in the consent process by demonstrating value in doing so,
- Be *robust to habituation and desensitization* over time,
- *Effectively present information* in a way that aids comprehension, and
- *Enhance multiple scales of use* without adversely affecting the overall informed consent process.

By "multiple scales of use" we refer to the natural variation in time and effort that users invest in the overall consent process. For example, some users will be more meticulous, while others will always want to rush through the process as quickly as possible (what Good *et al.* refer to as "drive-by installers" [8]). Given this natural variability, the system should enhance the consent process across all scales of use, without optimizing one scale of use at the cost of another.

The goals listed above consider the consent process from the user's perspective. It is equally important to consider the *implementers* and *designers* of any such system. Accordingly, an ideal consent process is:

- *Reproducible*, or easily instantiated with low cost and effort, and
- *Adaptable*, or easily modified to fit new contexts and needs

A consent process is easily *reproducible* if it can act as a *template* and be applied to new software systems with little modification. A process is *adaptable* if the *concepts* upon which it is based can be applied to significantly different contexts. For example, one might imagine a highly structured template for medical consent forms that is easily reproduced but not adaptable to new contexts such as software license agreements.

In considering all of these goals, it should be clear that they are often in conflict with one another. For example, being easily *reproducible* is at odds with *robustness to desensitization* because repeated exposure to the same style of informed consent can naturally desensitize a person to the process. Nonetheless, the goals represent ideals to strive for in the design of new consent processes.

With these goals in place, we now describe how we explored this design space.

Design Process

Noting people's tendency to skip text-based agreements, we started by considering a wide range of techniques to capture people's attention and to communicate the concepts found in consent agreements. These techniques ranged from supplementary videos and illustrations, to enhancements of the text-based agreement itself. We explored this design space by drawing inspiration from related fields (e.g., advertising, technical communications, comics) and engaging in rapid, low-fidelity prototyping and formative evaluations. Overall, 21 subjects participated in this phase of the research, evaluating dozens of prototypes.

Initial Lessons Learned

In evaluating the many prototypes, subjects often noted that redesigned consent processes gave the impression that the software distributors were *making an effort to communicate*. That is, subjects recognized an attempt to make the consent process accessible and meaningful, in contrast to current consent processes which merely appear perfunctory. This finding was encouraging, because it indicated the participants were interested in engaging in the consent process if it was apparent their best interests were in mind.

One of the more important lessons learned in this exploratory phase was that borrowing techniques from other forms of communication (e.g., advertisements) must be done with care to avoid causing users to confuse the consent process with these other forms of communication. For example, in one design we used a table to present information, rather than a bulleted list. The use of this technique reminded one user of technical manuals (specifically, troubleshooting tables), which confused the true intent of the document. Similar problems arose with video, where users assumed it was a training video, leading them to skip over it. Thus, while it can be fruitful to mix and match techniques from related areas, the final product must still clearly be recognizable as an intent to gain informed consent.

The third primary finding from this study was that subjects found the plain-text style of consent agreements to be off-putting, but responded positively to redesigned agreements that made more effective use of typography and graphic design. In considering the goals enumerated above, this finding was promising. In particular, we felt that there was the potential to cultivate a set of visual design techniques for improving consent agreements that could be easily reproduced and adapted with minimal cost and effort, compared to alternatives such as the use of video or custom illustrations. This inspired us to more methodically explore the range of ways to enhance the visual design of the consent agreement, which we describe next.

Gossip Magazines and Checkout Lines

To build a repertoire of visual design techniques to redesign consent agreements, we identified magazines as a potential source of inspiration. Magazine editors are highly motivated to create visual presentations that capture and retain individuals' attention [14] to increase sales, making them a rich source to draw upon.

We analyzed a range of magazines, including weekly celebrity magazines (e.g., *US Weekly*), women's magazines (e.g., *Glamour*), men's magazines (e.g., *Esquire*), news magazines (e.g., *Time*), and *National Geographic*. Two researchers went through each magazine page-by-page, circling and annotating techniques used to capture and retain attention.

Representative examples of each technique were torn out of the magazines and posted on a wall, where they were then clustered to form a large affinity diagram. These clusters led to a set of candidate design techniques that were iteratively developed, tested, and tuned to form textured agreements.

TEXTURED CONSENT AGREEMENTS

Textured agreements are visually enhanced consent agreements that employ typography, graphic design, and spatial layout to highlight information in the agreement and suggest its personal relevance. More specifically, textured agreements:

- *Layer* the presentation of information to make it easier to absorb the content at multiple scales of use,
- Employ *vignettes* that use mini-narratives to convey information in accessible formats,
- Use *sensationalism* to draw attention to particularly sensitive parts of the agreement, and
- Create *visual variety* in the agreement to counter lack of variety in the underlying text

While these methods are commonly used in other media, our contributions lie in the *selection* and *adaptation* of these strategies to the design of software consent agreements. Later, we will present experimental evidence supporting the effectiveness of these techniques. First, we describe each.

Information Layering

Information layering creates a *visual* hierarchy of the consent agreement's information, yielding a recognizable visual structure, as well as a clear path to navigating that information (Figure 1). [13,14,15] When done well, information layering helps capture a reader's interest and retain that interest, while assisting in the comprehension of the information presented. Furthermore, it accomplishes these goals at multiple scales of use.

Textured agreements layer information using a host of techniques, including "factoids" (described below), iconic symbols, warning boxes, and common conventions such as headings, bold text, and bulleted lists. These techniques (e.g., headings) are commonly used in other media. For the purposes of this paper, we explain the heuristics we developed for applying factoids, iconic symbols, and bold to consent agreements, as well as the intended effect of each strategy.

Factoids

In our study of magazines, we found it a common practice to include large *pull quotes* surrounding the main narrative. Pull quotes are catchy or interesting quotes taken from the primary text. Even after examining magazines for hours, we found ourselves continually drawn to these quotes, suggesting their robustness to desensitization. We generalize this convention to create *factoids* (Figure 2). Factoids are laid out similarly to pull quotes, but may contain content not within the actual consent agreement. Like pull quotes, textured agreements use factoids to pull readers in and convey key concepts through quickly consumed, "bursty" nuggets of information.

Factoids also help lend a human, personal element to the consent agreements both in the *form* and *content* of the factoids. More specifically, the act of singling out information and presenting it separately from the main text suggests that it is personally relevant and important that the reader notice it. Such devices also suggest that the document is *accessible* and that it includes interesting information, rather than just boilerplate legalese. For example, a factoid prominently placed at the top of one of our textured agreements exclaims that "less than 2% of the population reads these things," using humor to simultaneously acknowledge the tedium of reading a consent agreement, while suggesting that this consent agreement is different, accessible, and worth reading.

Iconic Symbols

Operational manuals often contain warnings to alert readers to information vital to their personal safety. These are frequently accompanied with an icon to demarcate the warnings (Figure 3). Textured agreements use warning symbols and colored boxes to highlight particularly sensitive information in the consent agreement, with the hope that even those who quickly skim the document will stop and read the content associated with the warning. Our formative evaluations suggest that these warning symbols must be used sparingly to ensure they are perceived as directing the user's attention to truly exceptional "hazardous" features of the agreement. Otherwise, users become desensitized to them and read only the first few warnings they encounter.

Heuristics for the Use of Bold

The use of bold in textured agreements is driven by holistic document needs as much as local, phrase-level needs. In particular, bold is applied to phrases across the document so that one can read only the bold phrases to gain a summary of the document.

Vignettes

To help draw users' attention, suggest personal relevance, and communicate vital concepts, textured agreements employ *vignettes*. Vignettes are mini-narratives related to the content of the agreement. For example, in some instantiations of textured agreements, a comic character is shown exclaiming, "This software collects WHAT?" suggesting the software collects personal information that some may find objectionable (Figure 4). The informal nature of the illustration and its suggestion of an underlying narrative adds interest to the content to pull the reader in and compels them to get engaged with both the narrative and the surrounding text.

Sensationalism

Many of the celebrity magazines reviewed (e.g., *US Weekly*, *Star*) use sensationalist rhetoric as a way to capture user attention. This device can be adapted to function in consent agreements as well, and is not difficult to achieve given the strict terms found in many agreements.

Textured agreements use sensationalist visual devices to point out potentially alarming aspects of the agreement. Factoids, comics, and section headlines can all be employed to create a sensationalist effect (Figures 4, 5). As much as

possible, these sensationalist elements highlight information likely to be personally relevant, and use methods such as comic characters put in situations that the user can relate to (Figure 5).

Visual Variety

Often there exist multiple ways to draw attention to elements in an agreement. For example, one could use headlines, warning symbols, or vignettes to highlight privacy information in an agreement. Textured agreements use *visual variety* to continually create interest in an otherwise lengthy document.

At a small scale, variety is introduced by manipulating typography (e.g., bold text, italics, and different font styles for body text and headings). At a larger scale, textured agreements use *progressive exposure to new design elements* to introduce variety. Visual elements, such as factoids, warning icons, and bullet points, are continually introduced throughout the agreements to entice readers to take notice.

Devices like factoids and vignettes are particularly effective means for introducing variety since they exist independent of the primary body of text and can thus be placed anywhere. Use of different colored backgrounds throughout the document (Figures 3, 4) can also help lend variety.

Side Effects of the Techniques

The application of these techniques creates a number of desirable side effects that can be leveraged to further engage readers. These include *attention hotspots* and *pacing*, which we describe next.

Attention Hotspots

The use of information layering techniques (e.g., factoids) and vignettes create *attention hotspots*, or areas of the document that stick out and compel users to read them. Textured agreements are designed to maximally use these hotspots. In particular, textured agreements strategically place hotspots near important content to increase the chance the nearby text is read (e.g., Figures 2, 4). We call this practice *attention piggybacking*.

Pacing

When principles such as information layering and variety are applied effectively, they help suggest an accessible *pacing* of the document. Pacing refers to a reader's assessment of the effort required to read a document [15]. Plain-text agreements create an impression that the document is lengthy and arduous to read. In contrast, textured agreements use strategies such as information layering to create a *textual pattern* that suggests that one can easily move through the document, at various levels of detail, to glean the information most relevant and important to them.

Summarizing the Techniques

Textured agreements visually enhance the presentation of consent agreements to improve their communicative effectiveness. Importantly, a textured agreement is not an immutable template. Instead, it represents a set of *strategies* one applies to the visual design of a consent agreement. These

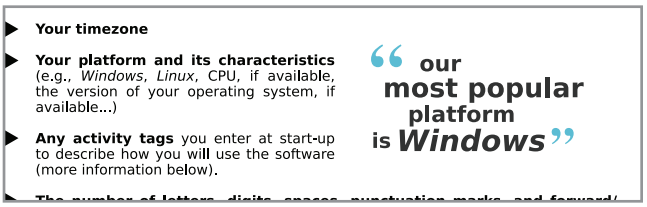

Figure 2. A factoid presenting information related to the nearby content in the consent agreement

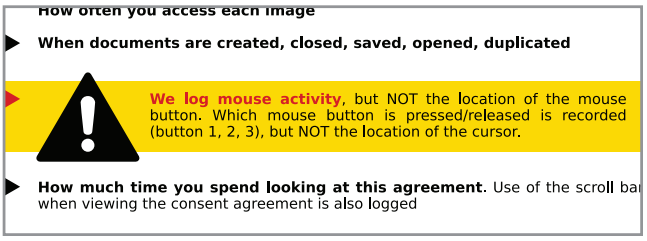

Figure 3. A warning box demarcated by an iconic exclamation point symbol

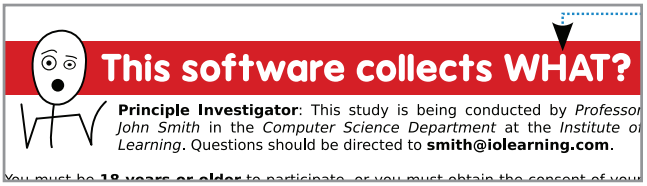

Figure 4. A sensationalist headline

strategies are *adapted* to each consent agreement by instantiating the elements in ways specific to the particular agreement. Notably, the selection of these strategies was intended to create a recognizable visual style that nonetheless incorporates variety (and hence, interest) each time it is instantiated. For example, factoids are a stable feature of textured agreements, but their content naturally varies from agreement to agreement to make each unique and interesting.

EXPERIMENTAL STUDY

Experimental Design

A between-subjects deception experiment was devised to evaluate textured agreements. The study employed five conditions corresponding to five different agreement styles: three conditions with textured agreements' visual techniques applied to varying degrees (minimal, moderate, heavy), a pre-installation summary condition similar to that of Good *et al.*'s study [7], and a control condition with plain-text consent agreements. The summary condition was included to partially replicate Good *et al.*'s previous study and provide another point of comparison for the textured agreements.

Subjects were asked to download, install, and use three image manipulation applications from a mock web page, for the purpose of choosing the best application for rotating images in a digital photo collection. Each application's installer was a custom-written installation program that could be experimentally manipulated to show a different consent agreement.

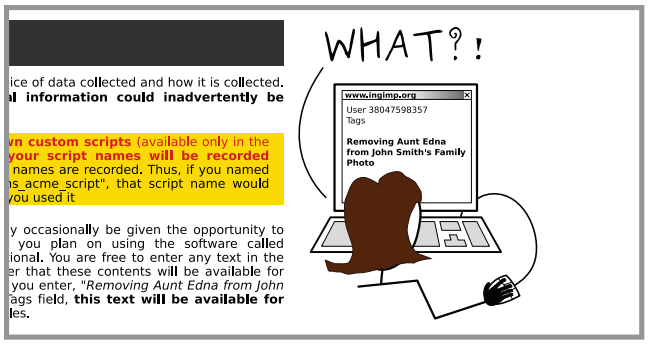

Figure 5. Vignettes draw the reader and communicate agreement content through mini-narratives

The installer was instrumented to collect timing information for time spent on each screen of the installer. After choosing an application, participants completed a questionnaire that included a content quiz to test how much information they absorbed from the consent process.

To increase ecological validity, we carefully designed the study to minimize the chance that participants would artificially focus on the consent agreements. We took two measures to avoid this potential bias. First, deception was employed. Subjects were told that the purpose of the experiment was to learn how they choose software when multiple choices exist. Second, a verbal consent script was used to obtain initial consent to participate in the study. This was motivated by the observation that written consent agreements primed users to look at the software agreements in formative evaluations. Subjects were debriefed upon completion about the true nature of the experiment and given a second, paper-based consent form to provide consent to keep their data.

Procedure

Subjects were given a written scenario and instructions after obtaining verbal consent. The scenario indicated that they had recently received a digital camera, but lacked software to perform basic manipulations of the images. Accordingly, they were told to imagine they had just found a website with three different photo applications. The instructions indicated that they could download, install, and use any of the applications. Their specific task was to decide which software they would choose to use for the purpose of rotating images. A folder of improperly-oriented photographs was provided to assist with their evaluations. Once they had made a choice, the instructions indicated that they should start a questionnaire to record their final selection.

After receiving and reading the instructions, subjects had the opportunity to ask questions. They were then seated at a desktop computer with a web browser already opened to the download web page to perform the task. After completing the task and the subsequent questionnaire, they were debriefed about the actual intent of the study and given a written consent agreement.

Experimental Conditions

We developed three separate plain-text consent agreements for the applications, drawing from existing software agreements. The content was designed to be consistent in form and presentation across all three consent agreements. Each agreement indicated that the application was instrumented to collect data, though the specifics of what was collected, why, and by whom, varied per consent agreement. A content quiz in the questionnaire asked which applications collected which data.

Three textured agreement *templates* were developed representing minimal, moderate, and heavy application of the techniques. These templates were applied to each of the plain-text agreements, yielding three separate instantiations of a template per condition (Figure 6 gives examples of each condition). We describe the features of the templates in more detail.

Based on experiences in the formative evaluation, we hypothesized that it was not sufficient to merely decorate a consent agreement to improve reading times. Instead, our experiences suggested that one needed to conscientiously apply the textured agreements' techniques. To test this hypothesis, the minimal condition represents an aesthetically pleasing consent agreement, but the design does not otherwise strive to reinforce the content of the agreement. The moderate condition represents what we feel is a balanced application of the strategies of textured agreements. This agreement uses a more professional font, warning boxes for three important agreement clauses, and factoids relevant to agreement content. Finally, the heavy condition was designed to incorporate as many of the techniques as possible to make a visually dense, deeply layered consent agreement (Figures 1 and 6C show examples of this agreement).

One-page summaries for the three consent agreements were developed using the heuristics used by Good *et al.* [7,8] (Figure 6D).

To avoid order effects, the order of agreements paired to installers was varied. Additionally, the order of application names on the web page was varied. To avoid potential effects due to pairing the same application names to the same agreements, there were six different application name-agreement pairings. This yielded a 5x6 between-subjects design with 5 conditions and 6 application name-agreement orderings.

Performance Measures

We measure the effect of agreements by measuring the amount of time spent on individual installer screens and the maximum amount the consent agreement is scrolled.

Because we strove to create a more realistic experience, users could run each installer multiple times. However, this capability complicates measures of time spent in the installer screens. Accordingly, our timing measures are derived from the *first* time an installer is run, which we refer to as *first-run timings*. This measure is likely to more closely correspond to actual practices since people typically only run an installer once. For each subject, an average timing is calculated from first runs of the installers. If an installer was never run by a subject, it is not included in the calculation of the average time. The maximum amount a consent agreement is scrolled is calculated the same way.

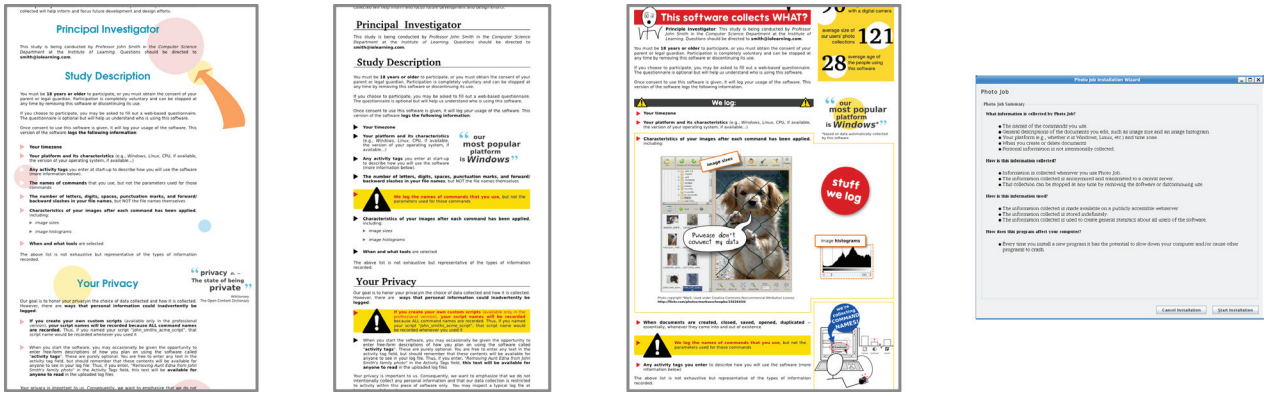

A. Minimal **B. Moderate** **C. Heavy** **D. Summary**

Figure 6. Examples of each experimental condition. For the textured consent agreements (A-C), the same section of each agreement is shown. For the summary condition (D), the initial summary screen is shown.

The effect of the techniques was also measured by the content quiz, mentioned earlier. Again, this quiz asked questions related to the software agreements' content (specifically, what data were and were not collected).

Experimental Apparatus

The study was conducted using a basic Windows XP installation on a VMWare virtual machine that had the Internet connection disabled. Using the snapshot feature of VMWare, we were able to have identical start conditions for every subject. A setup script was run before each session to set the experimental conditions for each subject.

The applications' installer was a custom-developed installer written in Java. It used the Windows "look-and-feel" to appear as a normal Windows application, and mimicked the appearance and feature set of a typical Windows installer. The installer was instrumented to record when it was opened and closed, the time spent on each screen, interactions with controls, and whether subjects canceled or completed the installation.

With the exception of the summary condition, each installation process was identical. The first screen was a welcome screen; the second, the consent screen; the third, a screen informing the user where the software will be installed; the fourth, a screen showing the installation progress; and the fifth, a screen indicating the installation was finished. The summary condition added an additional screen, a summary of the consent form, which was shown before the welcome screen. This ordering of screens mirrors that of Good *et al.*'s study, though we did not include a blank screen before the welcome screen in non-summary conditions. While they included this screen to act as a control for the presence of the summary screen, we wanted to increase the ecological validity of our study, and thus did not include it.

The installation program did not actually install the application (they were already installed), but did copy all of the files to a temporary directory to simulate the installation process. It also added shortcuts to the Start Menu.

Participants

90 subjects were recruited in a university setting. Six dropped out, providing 84 subjects, or 17 subjects per condition, with the exception of the heavy condition, which had only 16. Subjects were compensated with a $10 gift certificate for a coffee chain.

43 females and 41 males participated, aged 17-47 years old (mean=24, SD=6). Subjects' self-reported computer expertise on a five point scale was an average of 3.4 (SD=1) with 5 being "most expert."

RESULTS

Timings, Scrolling Behavior, and Reading Self-Reports

An analysis of consent screen timings indicates three outliers, one each in the minimal, summary, and control condition. These three subjects spent an average time of 250, 579, and 433 seconds, respectively, on the consent screen, each of which is more than 3 standard deviations from the within-condition mean. These outliers were dropped from the timing analyses and are not represented in any graphs presented here. Apart from these outliers, we observed considerable variation in reading habits. In the questionnaire, subjects were asked to self-report their tendency to read agreements. We found this measure of reading habits to be a contributing factor to the scores, and thus include it as a factor in our analyses.

Figure 7 presents a box plot of the first-run consent screen timings for each condition. Table 1 summarizes the data. The heavy condition features the longest consent screen time (mean=39.8 seconds, SD=39), followed by the moderate condition (mean=35.6, SD=39.2). Mean times in the summary and control conditions were 10.3 and 7.1 seconds, respectively. An ANOVA indicates significant differences between conditions ($F_{4,76}=5.65$, $p < 0.01$). Post-hoc Tukey tests indicate significant differences between: heavy and control ($p < 0.01$), heavy and summary ($p = 0.026$), moderate and control ($p = 0.026$), and a trend for significance between moderate and summary ($p = 0.063$). No other significant differences were found between conditions.

Measure	Heavy	Moderate	Minimal	Summary	Control
First-run consent screen times (seconds)	mean=39.8, SD=39.0	mean=35.6, SD=39.2	mean=16.7, SD=23.5	mean=10.3, SD=22.0	mean=7.1, SD=11.2
First-run maximum scroll amount (%)	mean=53.4, SD=48.3	mean=41.5, SD=38.5	mean=34.7, SD=38.7	mean=14.6, SD=29.1	mean=18.4, SD=32.0
Total time on pre-install screens (seconds)	mean=51.4, SD=49.2	mean=58.1, SD=56.7	mean=32.7, SD=40.0	mean=38.1, SD=38.9	mean=19.6, SD=22.5
Total maximum scroll amounts (%)	mean=54.2, SD=48.5	mean=53.3, SD=41.0	mean=43.1, SD=41.1	mean=19.8, SD=35.1	mean=34.1, SD=42.4

Table 1. Summary of results across conditions

An ANOVA indicates significant differences between conditions for scrolling, as well ($F_{4,76}$=3.96, p = 0.014). Post-hoc analysis indicates significant differences between heavy and control (p = 0.04), and heavy and summary (p = 0.02).

The questionnaire asked subjects to self-report how much they read each consent agreement. Self-reported reading amounts were found to be significant with respect to condition ($F_{4,76}$ = 3.16, p < 0.05), with post-hoc analysis indicating the differences are due to subjects reporting they read the consent agreements more in the moderate condition than in the control condition (p < 0.05).

Content Quiz Performance

Analysis of variance indicates no significant differences in performance on the content quiz ($F_{4,76}$=1.03, p=0.42). However, a positive correlation was observed between reading time and performance on the quiz (r = 0.56, p < 0.001), suggesting its validity as a measure.

Aggregated Timings Across All Installations

In examining the data, one thing we noticed was a tendency for subjects to cancel installations at the *summary* screen. This observation echoes Good *et al.*'s findings and suggests the effectiveness of the summaries in communicating information. However, this finding slightly complicates comparisons of time spent on the consent screen since subjects may not reach the consent screen the first time the installer is run. Accordingly, we define *total pre-install time* as the sum of time spent on all screens up to, and including, the actual consent screen. In the case of the summary condition, this includes the initial summary screen, the welcome screen, and the consent screen. For all other conditions, it includes only the latter two screens. For this measure, we also sum the time spent in these screens across *all* installation runs. While this measure is not perfect (since there is variability in how frequently people ran the second installer, and each run adds to this measure), it helps assess the potential impact of summaries by summing time across screens. We define *total, maximum scroll amounts* in the same way. Figures 8 and 9, and Table 1, summarize these measures.

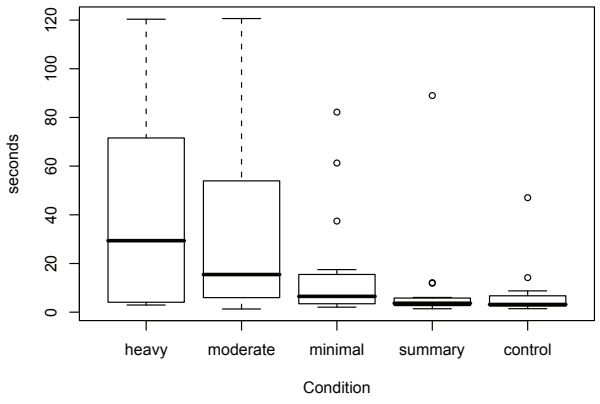

Figure 7. First-run time spent on consent screens

An ANOVA indicates no significant differences between treatment conditions for total pre-install time, though there is an apparent trend (p = 0.089). These trends appear to be due to both the moderate and heavy conditions. An analysis of scrolling behavior across all installation runs (rather than just the first run of each installer) supports this hypothesis. ANOVA of the scrolling behavior shows significance ($F_{4,76}$ = 3.36, p = 0.026), with post-hoc analysis indicating the differences are due to differences between the heavy and summary condition (p < 0.05) and the moderate and summary condition (p < 0.05), with subjects in the moderate and heavy conditions scrolling the consent agreements significantly more than in the summary condition.

Preferences

Subjects were asked to rate the visual appeal of the consent agreements. An ANOVA indicates a significant difference in visual appeal between conditions ($F_{4,76}$ = 7.61, p < 0.001) with post-hoc analysis revealing significant differences between heavy and control (p < 0.01); heavy and summary (p < 0.01); moderate and control (p < 0.05); and moderate and summary (p < 0.05). The minimal condition was also found more appealing than the summary condition (p < 0.05) and trended towards being more appealing than the control condition (p = 0.10).

Qualitative Feedback

In the questionnaire, subjects were asked to provide qualitative feedback regarding the visual appearance of the consent agreements and their overall informativeness. A number of comments suggest the techniques worked as intended. For example, a participant in the heavy condition commented:

> It got me to read them, when I install other programs, I NEVER read them. Big letters, organized points, and cartoons help. I think the organization was the most important.

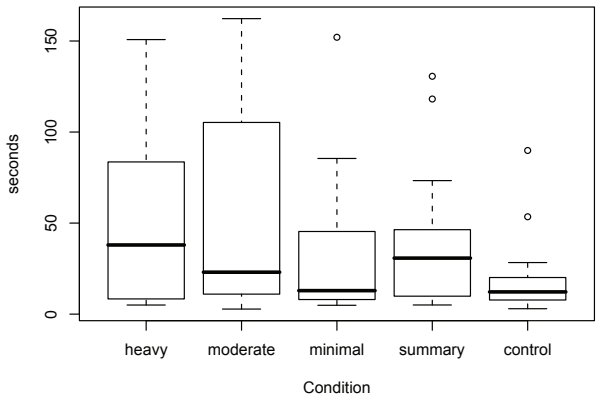

Figure 8. Total time on pre-install screens (seconds)

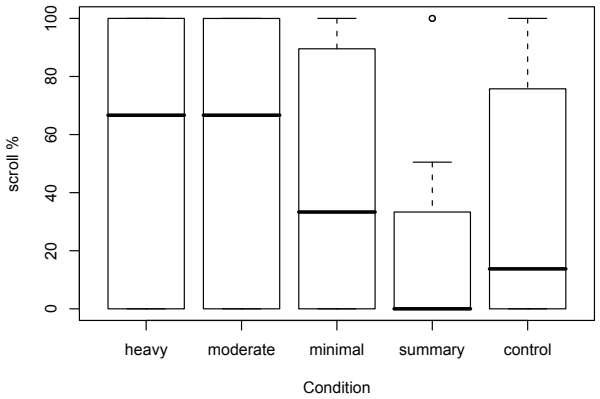

Figure 9. Total amount of scrolling of consent agreements, over all installation runs (%)

A participant in the moderate condition suggests the effectiveness of the documents' pacing:

> Compared to other agreements on other programs, this one is appealing because of the headings; which are placed similar to a newspaper to get one's attention. On other programs, it is just a bunch of words bunched together, similar to a contract but on the monitor.

Comments from subjects in the minimal condition indicate they noticed the changes in visual appearance. However, they did not comment on the organization of information or the notion of being compelled to read, as in the moderate and heavy conditions.

The comments suggest the techniques did not have universal appeal. For example, a subject in the heavy condition wrote:

> I think I would have read it more closely if it had been a little less over-the-top. I did really like that it was different and caught your attention.

Another subject noted that the heavy style was "somewhat obnoxious in coloration and layout." These comments suggest that while the heavy application of techniques attracted attention for some, it may be too much; the moderate application seems to strike the right balance.

DISCUSSION

The results of this experiment strongly support the notion that textured agreements compel people to read them compared to plain-text consent agreements. In particular, textured consent agreements increased the time spent on consent screens from an average of 7 seconds in the control condition to an average of 36-40 seconds in the moderately and heavily designed textured agreements. Notably, these effects were not observed in the minimal condition, suggesting this increase in time cannot be attributed to novelty alone.

The success of these techniques in increasing reading times is promising for two reasons. First, they represent only a subset of the wide repertoire of visual design techniques available; there is significant room for further improvement. Second, the technique achieved its effectiveness by operating on the *primary object of interest*, namely, the consent agreement itself. As we discuss next, there are good reasons to focus on improving this document, rather than introducing auxiliary documents.

Assessing Summary Condition

The summary condition was not shown to significantly affect the reading time of the actual consent agreements, whether one considers first-run times or total pre-install times. However, this finding does not indicate they are ineffective. A comparison of time spent on the summary screen compared to time spent on the welcome screen of the other conditions indicates that users *do* read the contents of the summary screen, spending 18 seconds longer on this screen compared to the welcome screen of the other conditions ($p < 0.0001$). However, our data indicate that there are side effects associated with reading the summary. In particular, the data argue that summaries reduce the likelihood that people spend time reading the *actual* consent agreement. Thus, while they effectively communicate a condensed version of the summary, they do so at the cost of reading the full agreement.

Human Ethics Perspective

Having observed the positive effects of the textured agreements, we met with three members of our internal review board who regularly review human ethics applications. We presented the textured agreements to gain their perspectives and understand potential issues in using them in practice.

The reaction to the agreements was extremely positive. Compelling study participants to read consent agreements is a problem they struggle with in study designs, so they welcomed the visual redesigns. However, they did have some suggestions for improving the designs and for future research. In particular, they observed that the heavily textured agreements could be problematic for seniors. This population might find the dense clustering of information distracting or difficult to comprehend. This point raises an important issue for future work: Examining potential age differences related to the particular designs. Apart from this suggestion, they felt the empirical evidence was reason enough to allow use of these agreements for software-based human subjects research targeting a similar demographic to that of the study.

CONCLUSION AND FUTURE WORK

This paper has introduced textured agreements, visually redesigned consent agreements that draw upon strategies of popular media to gain users' attention, retain that attention, and highlight information of personal relevance. An empirical study suggests these techniques show promise in improving the software consent process. The study results also suggest caution in using techniques that partition the consent process into multiple phases, as is done with summaries of agreements. While summaries are effective at conveying a synopsis of the agreement, they can lead users to ignore the full agreement.

In the future, we would like to perform longer-term evaluations of these strategies to determine their robustness with respect to desensitization and habituation. Additionally, we wish to study different demographic groups, particularly seniors, to understand how these techniques affect readability and reading behaviors. Finally, we would like to apply these techniques to a number real-world agreements to better understand their reproducibility and adaptability.

REFERENCES

1. Abrams, M., M. P. Eisenhauer, and L.J. Sotto. Letter to Federal Trade Commission. March 29, 2004. Re: alternative forms of privacy notices, project no. P034815. Hunton & Williams: The Center for Information Policy Leadership.
2. Acquisti, A., and J. Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE Security & Privacy* 3(1): 26–33.
3. Brustoloni, J. C. and Villamarín-Salomón, R. 2007. Improving security decisions with polymorphic and audited dialogs. In *Proceedings of the 3rd Symposium on Usable Privacy and Security* (Pittsburgh, Pennsylvania, July 18-20, 2007). SOUPS '07, vol. 229. ACM Press, New York, NY, 76-85.
4. Campbell, F. A., Goldman, B. D., Boccia, M. L., & Skinner, M. (2004). The effect of format modifications and reading comprehension on recall of informed consent information by low-income parents: A comparison of print, video, and computer-based presentations. *Patient Education and Counseling*, 53(2), 205-216.
5. Friedman, B., Lin, P., & Miller, J. K. (2005). Informed consent by design. In L. Cranor and S. Garfinkel (Eds.), *Designing secure systems that people can use* (pp. 495-521). Cambridge, MA: O'Reilly and Associates.
6. Friedman, B., Smith, I. E., Kahn, P. H., Jr., Consolvo, S., & Selawski, J. (2006). Development of a privacy addendum for open source licenses: Value Sensitive Design in industry. *Proceedings of Ubicomp 2006* (pp. 194-211). Berlin, Heidelburg, New York: Springer-Verlag.
7. Good, N. S., Grossklags, J., Mulligan, D. K., and Konstan, J. A. 2007. Noticing notice: a large-scale experiment on the timing of software license agreements. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA, April 28 - May 03, 2007). CHI '07. ACM Press, New York, NY, 607-616.
8. Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D., and Konstan, J. 2005. Stopping spyware at the gate: a user study of privacy, notice and spyware. In *Proceedings of the 2005 Symposium on Usable Privacy and Security* (Pittsburgh, Pennsylvania, July 06 - 08, 2005). SOUPS '05, vol. 93. ACM, New York, NY, 43-52.
9. HIPAA Highlights Privacy Notice, Press Release, Center for Information Policy Leadership, Hunton and Williams. http://www.hunton.com/news/news.aspx?nws_pg=7&gen_H4ID=10102 (last accessed Sep 19, 2008)
10. Jensen, C. and Potts, C. 2004. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vienna, Austria, April 24 - 29, 2004). CHI '04. ACM Press, New York, NY, 471-478.
11. Notes from Interagency Public Workshop: Get Noticed: Effective Financial Privacy Notices, (Dec. 4 2001) at http://www.ftc.gov/bcp/workshops/glb/. (Last accessed Sep 19, 2008).
12. Spiekermann, S., J. Grossklags, and B. Berendt. 2001. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proceedings of the Third ACM Conference on Electronic Commerce*, Association for Computing Machinery (ACM EC'01), 38-47. New York: ACM Press.
13. Tufte, E. R. (1990). *Envisioning information*. Cheshire, Conn.: Graphics Press.
14. White, J. V. (1982). *Editing by design : A guide to effective word-and-picture communication for editors and designers* (2nd ed.). New York: Bowker.
15. Wright, P. (1977) Presenting technical information: a survey of research findings. *Instruction Science*, 6, 93–134.