

Informed Consent Online: A Conceptual Model and Design Principles

Batya Friedman

The Information School
University of Washington
Box 352840
Seattle, WA 98195 USA
+1 206 523 1794
batya@u.washington.edu

Edward Felten

Dept. of Computer Science
Princeton University
35 Olden Way
Princeton, NJ 08544 USA
+1 609 258 5906
felten@cs.princeton.edu

Lynette I. Millett

Dept. of Computer Science
Cornell University
Ithaca, NY 14853 USA
+1 202 334 3347
millett@cs.cornell.edu

ABSTRACT

We provide a conceptual model of informed consent online. This model is based on five components: disclosure, comprehension, voluntariness, competence, and agreement. We examine how these components play out in a wide range of online interactions. Moreover, we call attention to the critical role of Web browsers in this endeavor. Finally we offer eight design principles for realizing informed consent online. This work fits within the emerging field of Value-Sensitive Design.

Keywords

Informed consent, computer ethics, cookies, discussion groups, e-commerce, ethics, human-computer interaction, human values, interface design, Internet Explorer, locus of control, Netscape Navigator, online interactions, online chat rooms, privacy, security, recommendation systems, social computing, social impact, tracking, Value-Sensitive Design, Web browsers, World Wide Web.

INTRODUCTION

Informed consent provides a critical protection for privacy, and supports other human values such as autonomy and trust. Yet currently there is a mismatch between industry practice and the public's interest. According to a recent report from the Federal Trade Commission [9], for example, 59% of sites that collect personal identifying information neither inform Internet users that they are collecting such information nor seek the user's consent. Yet, according to a Harris poll [2], 88% of users want sites to garner their consent in such situations.

The Federal Trade Commission hopes that industry will continue to make progress on this problem, in conjunction with its proposed legislation [9, p. iv]. Toward such progress, however, we in the HCI community should be helping to shape the dialogue. Specifically, we need to be seeking a clearer understanding of what constitutes informed consent, and of how it can be realized in online interactions.

Accordingly, in this paper we first provide a conceptual model of informed consent online. Our analyses here focus in particular on the relationship between users and Web sites. We employ illustrative examples from online recommendation systems, e-commerce, interface design, discussion groups, and chat rooms. Next we examine how technology – the Web browser in particular – plays a critical role in mediating between users and Web sites. Finally we propose eight design principles for realizing informed consent in online interactions.

WHAT IS MEANT BY INFORMED CONSENT?

Informed consent – both words carry import. Others have suggested [1, 3], and we concur, that the idea of “informed” encompasses disclosure and comprehension. In turn, the idea of consent encompasses voluntariness, comprehension, and agreement. Thus to understand the general idea of informed consent, we first need to focus on each of these five conceptual components. We do so now, paying particular attention to how they apply in a wide range of online interactions.

Disclosure

Disclosure refers to providing accurate information about the benefits and harms that might reasonably

be expected from the action under consideration. What is disclosed should address the important values, needs, and interests of the individual, explicitly state the purpose or reason for undertaking the action, and avoid unnecessary technical detail. The information should also disabuse the individual of any commonly held false beliefs. Moreover, if the action involves collecting information about an individual then the following should also be made explicit: (a) what information will be collected; (b) who will have access to the information; (c) how long the information will be archived; (d) what the information will be used for; and (e) how the identity of the individual will be protected.

Consider, for example, online interactions with an e-business that builds and maintains a recommendation system based on customers' previous purchases. The e-business should disclose to the customer the purpose and benefits of the recommendation system (e.g., to aid future customers with future purchases based on information gleaned from previous customer purchases) as well as potential harms (e.g., in certain circumstances it may be possible to identify the user). The e-business should also inform customers about what information they will retain (e.g., the customer's name, what was purchased), who will have access to that information (e.g., other customers using the recommendation system, programmers, any third party companies), how long the information will be archived (e.g., one year? Indefinitely?), what the information will be used for (e.g., building the recommendation system, other uses?); and how the identity of the individual will be protected (e.g., no identities will be revealed explicitly through the recommendation system, efforts will be taken to remove other identifying information such as geographic location).

Comprehension

Comprehension refers to the individual's accurate interpretation of *what* is being disclosed. This component raises the question: What criteria must be satisfied in order to say that something has been adequately comprehended? While there is no easy answer here, at least two methods seem viable: (1) being able to restate in different words what has been disclosed, and (2) being able to apply what has been disclosed to a set of hypothetical events. For example, continuing with the recommendation system example above, based on what has been disclosed can the user answer reasonable questions about the data's use such as: Will information

about the user's last three prior purchases be included in the recommendation system? Will someone using the recommendation system be able to determine what the user has purchased in the past? Will information about the user's past purchases be a part of the recommendation system two years from now?

In face-to-face interactions – a common means to obtain informed consent in the research, medical and other communities – the individual can ask questions and a professional can engage the individual in further dialog (e.g., countering incorrect beliefs specific to that individual). Such dialog, as well as facial and other physical cues, help ensure there has been adequate interpretation of the information disclosed. Online interactions, however, lack many of the opportunities in face-to-face interactions to ensure and ascertain comprehension. In typical Web-based interactions, users are presented with a Web page or dialog box containing disclosure information and provided with an opportunity to agree or decline to participate by clicking on a button. Rarely are email or chat facilities used to provide even rudimentary opportunities for further dialog.

Granted, it is not possible to guarantee comprehension for all individuals in all situations. However, special efforts on the part of Web designers will be needed to analyze what users need to understand about a particular disclosure and to utilize strategies to increase the likelihood that comprehension will be realized.

Voluntariness

With voluntary action an individual could reasonably resist participation should he or she wish to. Voluntariness, then, refers to ensuring that the action is not controlled or coerced.

Coercion is an extreme form of influence that controls by compulsion, threat, or prevention and thereby violates the component of voluntariness. The canonical example of coercion occurs as follows: Person A holds a gun to Person B's head and says, "Fly me to Havana or I'll shoot."

A less obvious form of coercion can occur when there is only one reasonable way for individuals to receive certain needed services or information (or, if other ways do exist, those ways are too costly in terms of finance, time, expertise or other costs to be viable options). For example, currently in order to be placed in a residency program newly matriculated medical students must participate in

the National Medical Residency Program (NRMP) that matches residency openings (hospitals) with students. Medical students who elect not to participate in the NRMP are automatically excluded from virtually all of the residency openings at all of the best teaching hospitals [7]. Thus, medical students who desire a placement at the best teaching hospitals in effect have no option but to participate in the NRMP.

This less obvious form of coercion is a serious concern for online interactions. Currently (and for the foreseeable future), Web sites engage in strikingly similar practices with respect to informed consent. The same can be said for the features and services provided by Web browsers. If there is a service or information that individuals need to obtain online – as will increasingly become the case as job advertisements, applying for medical insurance, applying for admittance to higher education, and other critical services move online in their entirety – then to participate in these services individuals must engage in Web interactions. Given the lack of substantive choice among Web sites and Web browsers, users can be in effect coerced into Web-based interactions that compel them to give up personal information or engage in other activities.

Manipulation of certain forms can also undermine voluntariness. Manipulation can roughly be defined as “any intentional and successful influence by a person by noncoercively altering the actual choices available to the person or by non-persuasively altering the person’s perceptions of those choices.” [3, p.354]. The key here is that manipulation alters the individuals’ choices or perception of choices by some means other than reason.

Manipulation of the sort we are concerned with can be achieved in at least three ways. One way entails manipulation of the options presented to the individual such that the presentation encourages certain choices or behaviors. For example, consider an e-business that asks the user for more information than is necessary to complete a purchase but does not indicate to the user that completing some fields is optional. Here, the user has more options than the e-business has made clear and the user has no way of knowing otherwise.

A second way entails manipulation of information. This manipulation uses information intentionally to overwhelm the individual or to provoke or take

advantage of an individual’s fear or anxiety. For example, in recent years some Web sites have packaged information into multiple cookies – information that could have been packaged more concisely – so that the user who elects to agree to each individual cookie is bombarded with numerous requests to set cookies from a single site. This overwhelming request for information to accept each cookie has at least two effects: (1) to influence the user to turn off the “agree to each cookie” option in order to avoid the overwhelming request for information, and (2) to increase the likelihood that the user would fail to notice a cookie that the user might wish to avoid, such as a third party cookie intermingled among the many cookies from the target Web site.

The third form of manipulation is psychological. This form of manipulation includes any intentional act that influences a person by causing changes in the individual’s mental processes by any means other than reason. Flattery, guilt-induction, and subliminal suggestions are a few relevant influences. Recent work by Reeves and Nass and their colleagues [4, 5] indicates that individuals are vulnerable to psychological manipulation in online interactions, particularly with respect to psychological manipulations from the technology’s interface. For example, in their research, Reeves and Nass have shown that users respond to flattery from a computer, judge computers that criticize rather than praise others to provide more accurate information, and apply gender stereotypes to the technology simply on the basis of subtle gender cues, to name but a few of their results. Web sites that use psychological manipulation – for instance, to flatter the user into divulging information or into attributing greater accuracy to online recommendations -- may violate the criterion of voluntariness.

Competence

Competence refers to possessing the mental, emotional and physical capabilities needed to be capable of giving informed consent. For example, a person with Alzheimer’s may lack the mental capability to make his or her own medical decisions. Or, in the online environment, a 15-year-old may lack the mental and emotional capability to make reasoned judgments about when to provide personal information to e-businesses and in online chat rooms.

Young people may be particularly vulnerable online as their technical competence typically far exceeds their mental and emotional development.

For example, at roughly the same time as the Year 2000 Census was being conducted in the United States, the Barbi Web site presented Barbi as a census taker who asked Web site visitors -- mostly young girls under the age of twelve -- to help Barbi with her census work by completing a form that requested personal information. Troublesome from the perspective of informed consent, these young girls often lacked the mental and emotional competence necessary to judge the appropriateness of the information they volunteered to the Web site.

Web designers of sites targeted for children and adolescents will need to be especially cognizant of the component of competence. Drawing from the dictates of academic research Institutional Review Boards, we offer the following rule of thumb: For children younger than eight years of age, obtain informed consent from the child's guardian; for children eight years of age through (and including) seventeen, obtain informed consent from both the young person and the young person's guardian. Note that the challenges of comprehension are likely different for the young person and the guardian.

Agreement

Agreement refers to a reasonably clear opportunity to accept or decline to participate. In online interactions, opportunities to accept or decline should be visible and readily accessible. Opportunities to accept or decline that are buried under layers of menus or hidden in obscure locations are at best marginally viable.

In traditional human subjects research, the component of agreement is on-going. Participants may withdraw their agreement to participate at any time and for any reason (indeed, participants do not need to provide a reason for discontinuing participation). While the arena of research differs in important ways from online interactions and considerable complexity exists concerning how to apply the guidelines from one to the other, still the aspect of on-going agreement may have relevance for online interactions. For example, in the case of recommendation systems users could be provided with the opportunity to withdraw their data from the recommendation system at any time. Or with cookies, users could be provided with an easy mechanism to delete a cookie or to change a cookie's expiration date.

A related issue for on-going agreement arises in the context of discussion groups and online chat rooms

where dialog that may often feel like "ethereal" online conversation is archived, and in reality is more permanent and accessible than most other forms of communication. In these online forums, participants in the flurry of heated conversation may forget they have agreed to have their online conversation recorded and archived and, if given the opportunity, would suspend their agreement at that moment in time. Mechanisms that periodically reminded participants that online dialog was archived (and perhaps allowed participants to remove dialog from the archive) could help preserve informed consent in these interactions.

Finally, not all forms of agreement need be explicit. As a society, we have a good deal of experience with implicit consent where by virtue of entering into a situation the individual has in effect agreed to the activities that are broadly known to occur in that context. For example, when a player steps out onto the football field in football garb and enters the game, the individual has implicitly agreed to participate in the normal activities of the game, namely, to being bumped, bashed, and smashed by other players who have entered into identical agreements. Implied consent holds in this case because the other components have also been met: disclosure and comprehension (via reasonable expectation), competence (if we assume the individual is of a reasonable age and of sound mind and body) and voluntariness (if we assume the individual was not coerced or manipulated to dress in football garb and to go out onto the field). For implied consent to hold for online interactions, similar criteria need to be met.

THE UNIQUE ROLE OF THE WEB BROWSER

We now call attention to how the Web browser mediates much of what occurs between the user and a target Web site and, thus, plays a critical role for informed consent in Web-based interactions. In its most general sense, browser software mediates communication between a client (typically an end-user) and a server (typically a remote machine serving a Web site). For purposes of simplicity, we will speak in terms of the user's computer and a Web site's server, though our discussion applies more generally to any client/server architecture.

The browser software occupies a special place in the user's interactive experience and, in effect, stands between the user and any remote connection. To achieve this functionality, the user's system runs the browser software and allocates to it memory and areas of the screen that are under the browser's control. When the user

asks to view a Web page, the browser sends a request to the appropriate Web server, and the server replies by sending a description of how to display the page. This description can be thought of as a series of requests to the user's browser, which the browser carries out in order to display the page. These requests are, of course, filtered through the user's Web browser - that is, the browser first receives and reviews the remote machine's request and then either fulfills or denies the request.

When a Web browser is explicitly programmed to fulfill a request from a remote site "to do something", we call that *enabling* a capability. That is, the Web browser is now capable of fulfilling a particular request¹. For example, most common browsers have been programmed with the capability to store cookies on the user's machine and to provide the remote site with a copy of the user's mouse movements. Correspondingly, when a Web site makes a request that a Web browser is capable of fulfilling, we call that *exercising* a capability.

After a remote site has exercised a capability, the Web browser software has no control over what the remote site does with the information or other actions that the site may take. It is worth noting that, in general, the user's browser cannot tell why a particular request was made, nor can it tell how a particular piece of information (such as the user's mouse position) will be used by the server once it has been made known to the server or to software sent by the server. The browser can act as a simple gatekeeper only, deciding which kinds of requests to allow and which kinds of information to reveal to the server.

In addition to determining what requests from the remote machine will be fulfilled, the Web browser plays at least two other critical roles with respect to informed consent online. First, the Web browser controls whether or not the user is notified about the request and, to a large extent, the content of that notification. In the case of a Web site's request

for the user's mouse position, the browser is positioned to inform the user about the request and, perhaps, to provide a structure that would encourage (if not require) the Web site to specify what the information will be used for. Thus, we see how the components of disclosure and comprehension largely reside with the Web browser. Second, the Web browser controls whether or not the user has an opportunity to agree to or decline the Web site's request. Thus the component of agreement also resides with the browser.

This is not to say that Web sites are entirely dependent on the Web browser for implementations of informed consent. A proactive Web site can implement its own version of disclosure and agreement that supplements that provided by the Web browser. But such implementations would be ad hoc and require users to become familiar with each Web site's policies and practices.

DESIGN PRINCIPLES FOR INFORMED CONSENT ONLINE

Based on all of the proceeding analyses, we now distill eight design principles for informed consent online.

1. *Decide whether the capability is exempt from informed consent.* Obtaining users' informed consent comes with a high cost to users. Information must first be disclosed and comprehended by users, who must then have an opportunity to agree or decline. And while all of this is being done, the user has been diverted from the task at hand - the thing that he or she really wanted to do. Let us refer to these costs for obtaining the user's informed consent as "the nuisance factor". If all Web-based interactions required explicit informed consent, the nuisance factor would be unmanageable. Online interactions would simply crumble under the burden.

Fortunately, a fair number of Web-based interactions may be considered exempt from the need to obtain users' informed consent.

But how are Web designers to determine which interactions are exempt? While there are no hard and fast rules, the Belmont Report [1] on human subjects research offers some useful guidelines. According to the report, an individual's participation is considered exempt from the need to obtain informed consent when the following three conditions are met: (1) Participation can in no way

¹ For completeness of discussion we note that the Web browser could be programmed to deny the request. There is also a third possibility in which the browser is not explicitly programmed to handle a particular request. In this instance the browser typically ignores the request (which has the same effect as to deny the request).

put the individual in physical, legal, psychological or social jeopardy. To this list of harms, for online interactions we also include that participation does not place the individual's privacy, data or hardware in jeopardy. (2) The purpose and sponsorship of the activity is known (or clearly stated to the individual). And (3) no coercion is involved. Designers can invoke these three criteria to scrutinize new Web-based interactions for exemption from informed consent.

The "canonical" case for exemption from informed consent entails a "safe" interaction that adds functionality for the user and Web sites but does not put the user, the user's data, or the user's hardware at risk. Moreover, the purpose of the interaction is obvious to the user and no-coercion is involved. The original HyperText Markup Language (HTML)² provides one example. Original HTML is comprised of a limited set of "safe" commands that control how the browser displays data (e.g., bold, centered) on the user's screen. Because the syntax of original HTML did not support activities that transmitted or changed the user's data (such as sending data to a remote machine or deleting data in the user's machine) no "unsafe" statements could be formulated. Thus, the very definition of the language ensured that no statement in the language could put the user or the user's data in jeopardy. In addition, the purpose of the HTML code was clear: to display data. And users were not forced to use HTML; they could obtain files in other formats. Thus, the use of original HTML meets the conditions for exemption. Web designers can comfortably make use of original HTML without obtaining the user's informed consent.

In an ideal world, most Web interactions as with original HTML would be clearly exempt from informed consent. But the world is far from ideal. For example, some Web interactions do not fully meet the criteria for exemption but, on balance, may not warrant the overhead of informed consent. In general, such cases occur when a Web interaction involves a minimal risk of harm or a minimal violation of the user's rights. Then the nuisance of obtaining informed consent may outweigh the potential harms. The nuisance factor gains further sway when we take into account the pragmatic concern that as the cumulative burden of obtaining informed consent increases, users may

² We limit our discussion here to the original HTML; dynamic HTML poses unique challenges for informed consent.

not attend to informed consent when it really matters. These cases will require careful judgment on the part of designers.

Granted it may be difficult to make these judgments in advance; however, defensible assessments should be made and, if the initial assessments are in error, then remedies should be implemented.

2. *Particular care must be taken when invoking the sanction of implicit consent for Web-based interactions.* On the surface implicit consent seems a reasonable umbrella to cover most Web interactions. After all, Web users do not interact online – as in the canonical example – with a gun held to their heads. However, unless users have comparable (with respect to costs such as time, effort, knowledge, and expense) alternative access to comparable services, products, information, and so forth, then Web use may not be regarded as wholly non-coercive. Given the rapidity and widespread movement with which access to goods and services have moved online and the corresponding movement to discontinue or minimize traditional means of access, the viability of alternative comparable access to goods and services is at times slim and getting slimmer. In this climate, we advocate presuming that implicit consent is not a viable option, and only in special circumstances and after careful consideration invoking the sanction of implicit consent.

Further challenges for implicit consent arise from the criterion of disclosure. The disclosure issue can be understood as follows: Although the user may be told what mechanisms are enabled, the user may not be aware of the full implications of those mechanisms. For example, while many users were aware of and enabled cookies, few users understood the implications of cookies for individual privacy (until an extensive public discussion took place).

3. *Defaults matter.* It is well established that most users do not change preferences settings. Thus, default settings should err on the side of preserving informed consent. Notably, the default setting for cookies on current browsers is "accept all cookies" which neither informs nor obtains consent from users. Thus casual users' "out-of-the-box" experience of cookies in mid-1999 was no different than their "out-of-box" experience of cookies in 1995. Default settings will also need to take into account the nuisance factor to obtain

users' informed consent (see Design Principle 4 that follows).

4. *Put users in control of the "nuisance factor".* Different users place differing degrees of importance on different types of harm. Correspondingly, how much of a nuisance factor a user is willing to tolerate to obtain informed consent will depend on the particular user. Rather than mandating a single mechanism for obtaining informed consent for all users in all situations, designers need to provide users with a range of mechanisms and levels of control so that users are positioned to manage the nuisance factor in accordance with their concerns. Successful designs will likely contain a reasonable balance among overarching controls (e.g., "never accept cookies" or "always accept cookies"), micro-managed controls (e.g., "ask about each cookie"), and intermediate controls that mix well-chosen overarching controls with micro-managed subsets (e.g., "decline all third party cookies and ask me about all other cookies").

5. *Avoid technical jargon.* Follow the well-established interface principle to avoid technical jargon in favor of clear language that directly addresses the user's values, needs and interests. For example, while technically precise the cookie dialog box text "The server so-and-so wishes to set a cookie that will be sent to any server in the domain X" would be better stated as "The Web site so-and-so wishes to set a cookie that will be returned to the Web site thus-and-such."

6. *Provide the user choices in terms of potential effects rather than in terms of technical mechanisms.* Even when technical jargon is avoided, users may be unaware of the implications – positive or negative – for selecting certain technical mechanisms. Whenever possible provide users with choices based on the type of protections a mechanism provides rather than by the name of the mechanism. For example, the preference setting label "Use RSA encryption" would be better written as "Encrypt message (to help prevent eavesdropping while your message is being delivered)."

7. *Field test to help ensure adequate comprehension and opportunities for agreement.* Because online interactions will likely rely on automated means to realize informed consent, Web designers face significant challenges in ensuring adequate disclosure, comprehension and opportunities for agreement. Thoughtful interface

design will need to be coupled with equally thoughtful field tests to validate and refine the initial designs. Moreover, because informed consent carries a moral imperative, the components of disclosure, comprehension and agreement need to work reasonably well for all users. Thus, it becomes a requirement (and not simply better practice) to include a reasonable range of both representative and atypical users in the field tests.

8. *Design proactively for informed consent.* More frequently than not, Web-based interactions and capabilities are conceived of and implemented without consideration of informed consent. Once introduced, online practices evolve around these new interactions and these, too, develop without consideration of informed consent. When issues of informed consent at last come to the fore, designers face a near insurmountable task: To retrofit the online capability and interaction. The solution, in part, is to design proactively for informed consent.

CONCLUSIONS

We have proposed a conceptual model of informed consent online. This model is based on five components: disclosure, comprehension, voluntariness, competence, and agreement. We examined how these components play out in a wide range of online interactions. Moreover, we called attention not only to the critical role of Web sites, but Web browsers, in this endeavor. Finally, we offered eight design principles for realizing informed consent online.

As with all design principles, these eight will need to be applied with good judgment. Moreover, at times two or more principles may come into conflict, and then a balance between what is gained and what is lost must be achieved. In other words, we are not arguing that informed consent can be fully realized in every situation. But drawing on the emerging field of Value-Sensitive Design [4, 5, 6], we are saying that Web sites and Web browsers, in addition to being judged by the quality of their content, presentation of information, features, and speed, should be judged on how well they substantiate critical human and particularly moral values. Informed consent is one such value that we hope will gain the eye of more researchers and designers in the field of Human-Computer Interaction.

ACKNOWLEDGMENTS

This research was funded by the National Science Foundation Award IIS-9911185. We thank Peter

H. Kahn, Jr. for his comments on earlier versions of the manuscript.

REFERENCES

1. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, Washington DC, 1978.
2. Business Week/Harris Poll. *A Growing Threat*.
http://www.businessweek.com/2000/00_12/b3673010.htm .
3. Faden, R, and Beauchamp, T. *A History and Theory of Informed Consent*. Oxford University Press, New York NY, 1986.
4. Friedman, B. (Ed.) *Human Values and the Design of Computer Technology*. Cambridge University Press, New York NY, 1997.
5. Friedman, B. Value-Sensitive Design: A Research Agenda for Information Technology. (Contract No: SBR-9729633). National Science Foundation, Arlington, VA, 1999.
6. Friedman, B., & Kahn, P. H., Jr. A Value-Sensitive Design approach to augmented reality. In W. Mackay (Ed.), *Design of Augmented Reality Environments*. MIT Press, Cambridge MA, in press.
7. Friedman, B., and Nissenbaum, H. Bias in computer systems. *ACM Transactions on Information Systems*, 14(3), 330-347, 1996.
8. Nass, C. I., Moon, Y., Morkes, J., Kim, E., & Fogg, B. J. Computers are social actors: A review of current research. In B. Friedman (ed.), *Human Values and the Design of Computer Technology* (pp. 137-162). Cambridge University Press, New York NY, 1997.
9. *Privacy Online: Fair Information Practices in the Electronic Marketplace* (A Report to Congress). Federal Trade Commission, Washington DC, May 2000.
10. Reeves, B., & Nass, C. (1996). *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, New York NY, 1996.