Privacy Critics: UI Components to Safeguard Users' Privacy

Mark S. Ackerman
Information and Computer Science
University of California, Irvine
Irvine, CA 92697 USA
ackerman@ics.uci.edu
http://www.ics.uci.edu/CORPS/ackerman.html

Lorrie Cranor
AT&T Labs-Research
Shannon Laboratory
Florham Park, NJ 07932
lorrie@research.att.com
http://www.research.att.com/~lorrie

ABSTRACT

Creating usable systems to protect online privacy is an inherently difficult problem. Privacy critics are semi-autonomous agents that help people protect their online privacy by offering suggestions and warnings. Two sample critics are presented.

KEYWORDS: privacy, World Wide Web, critics, agent architectures, CSCW, collaboration, P3P.

INTRODUCTION

Online privacy is a growing problem for Internet users. Of particular concern is the unanticipated release (and subsequent use or misuse) of personal information. As Goffman [5] noted, every individual wishes to present an appropriate "face" to the myriad of audiences: One may wish to be the dutiful worker to managers, but an unhappy employee to fellow union members. Everyday life requires that only the proper information be released at the proper time, and people do this seemingly without thinking about it. To lose control over this process is very disconcerting.

Currently Internet users have little knowledge about how information they release online will be used. Users who wish to engage in electronic commerce must often release personal information to complete transactions. However few web sites explain how that information will be used or whether it will be linked with other personal information [1], and in many countries (e.g., the US), few legal privacy protections exist.

Users would benefit from systems to assist them in identifying situations where their privacy might be at risk. However, as we shall explain, many aspects of privacy make it difficult to design usable systems. These usability issues have led us to construct *privacy critics*, agents that help users protect their privacy online. These critics currently work with the World Wide Web Consortium's Platform for Privacy Preferences Project (P3P).

PRIVACY AS AN INFORMATION INTERFACE PROBLEM

P3P is one attempt to address the desire for personal privacy along with the needs of electronic commerce. P3P allows Web sites to make statements ("proposals") about their privacy policies and request data using a standardized vocabulary and protocol [2]. Thus users will be able to make informed decisions about releasing personal information.

Unfortunately, P3P user interfaces suffer from a particular class of interface problem. The HCI restatement of the privacy problem reveals it to be wicked (in the computer science sense): The problem is inherently complex, ill-defined, and seemingly insolvable. This is true for not just one reason, but several.

If a person wishes to control what information she presents to whom, this results in an enormous information space (i.e. each datum a person has about herself against each person or organizational entity with which she comes into contact). Moreover, the space is actually more complex, since there are additional dimensions (e.g., what the organization wishes to do with the data, the degree of trust the individual has in the requesting entity). Clearly a matrix-style user interface for P3P over each of its ten dimensions would be overwhelming. On the other hand, simplified interfaces remove important detail for some users.

Furthermore, we noted above that an individual does not, in fact, deliberate within each social encounter. Therefore, the user's interaction with an interface for controlling private information must be nearly transparent and minimal during the actual social engagement.

Privacy, then, poses a very difficult HCI problem. Not only must a program present an extremely complex information and decision space, it must do so seamlessly and without interference in the natural progression of social engagements.

Simply put, we do not know how to design these kinds of interfaces. Yet, if this problem must be solved currently (and there are ample reasons to believe that it must be), then the resulting HCI challenge must be to find approximations for the problem that provide sufficient functionality as well as ameliorations to the secondary problems that will naturally occur from using approximations. The following discussion introduces privacy critics, semi-autonomous agents that help users protect their private information. We believe that privacy critics are both approximations and ameliorations to the privacy problem.

PRIVACY CRITICS

Critic-based architectures were first introduced by Fischer [3]. A critic, a type of intelligent agent, provides feedback and suggestions as users go about their ordinary tasks. For example, the HYDRA critics [4] provided design feedback for kitchen architects as they laid out kitchens.

Two important features of critics should be noted. First, they provide feedback to users - they do not necessarily take action on their own. This is an important distinction from other types of intelligent agents. Privacy critics, then, would help (rather than attempt to automate) the user's control over private information. They might offer suggestions or warnings to users, watching over their shoulders in a manner of speaking.

Second, a critic-based environment might have hundreds of different critics. Each would check on a different facet of a problem domain and user goal. There need not be (and usually will not be) one "true" privacy critic. The independent nature of the numerous critics allows one to consider an ecology of critics (to be discussed further below). Users are, of course, free to turn these critics off and on, set threshold levels, and decide what aspects of privacy they wish to guard most closely.

SAMPLE CRITICS

Privacy critics, then, are agents that watch the user's actions and make privacy suggestions. We have implemented prototypes of six sample critics; two are presented here. These six are merely the beginning of what can be done.

The first critic checks the simulated CyberPrivacy Advocacy Group’s database for consumer complaints about a Web site. We imagine a number of third-party databases collecting claims or problems about different kinds of sites. For example, a Better Business Bureau database could report that sites have had privacy complaints against them; other databases might report sites participating in data scams. This critic does not currently learn to categorize sites or learn about user preferences; these would be potential extensions.

The second critic watches the type of information being released and warns users when a P3P proposal requests data elements that can be used in combination to identify the user. For example, many people do not know that specific demographic data (e.g., race, birth date) can be used with zip code to uniquely identify individuals or households.

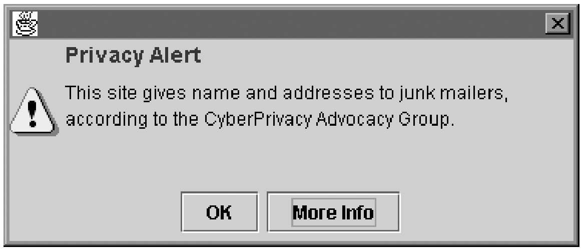

Privacy Alert

This site gives name and addresses to junk mailers, according to the CyberPrivacy Advocacy Group.

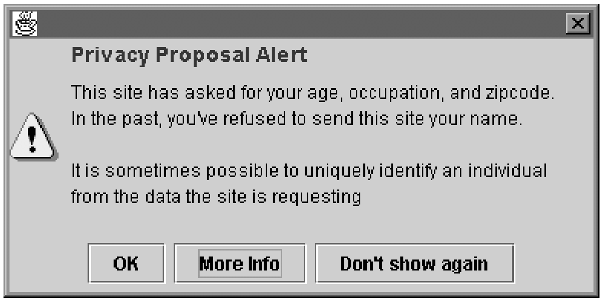

Privacy Proposal Alert

This site has asked for your age, occupation, and zipcode. In the past, you've refused to send this site your name.

It is sometimes possible to uniquely identify an individual from the data the site is requesting

IMPLEMENTATION AND FUTURE WORK

The construction of these critics, if they are to be viable, must occur at two levels. In addition to the critics themselves, a critic-based architecture must be implemented.

The current implementation of the sample privacy critics uses client-side proxies for prototyping. These proxies either intercept HTTP requests for URLs and simulate going to a third-party verifier, or they intercept simulated P3P proposals and make decisions on behalf of the user. (In P3P parlance, the proxy serves as a P3P user agent, incorporating a rudimentary trust engine to decide which proposals should be accepted.) Each critic has been separately implemented, using Java. While limited, these initial prototypes have been valuable for informal user testing. Feedback from users (college students) indicates that the idea of a privacy critic is relatively straightforward to explain and understand, and that once understood, the idea is even exciting to users.

The second level of implementation is a general user agent architecture that allows a range of critics. In order to have a flourishing ecology of privacy critics, third parties must be able to create new critics. As mentioned, we would like users to be able to add or remove critics, and to be able to obtain new critics as situations demand. For example, as new information scams spread across the Internet, it will be important to obtain the latest critics. Vendors of browsers may provide user agents with limited protection for users; users could then obtain additional privacy critics from consumer advocacy groups, trusted third parties, small companies, or hobbyists. This ecology of critics can occur, however, only if the architecture for the P3P user agent is suitably open. We are currently designing the necessary support services for such an architecture.

REFERENCES

1. Cranor, L. Internet privacy, a public concern. *netWorker: The Craft of Network Computing*, June/July 1998, 13-18.
2. Cranor, L. and J. Reagle. The Platform for Privacy Preferences. *Commun. ACM*, 42(2), in press.
3. Fischer, G., A. Lemke, T. Mastaglio and A. Morch. Using Critics to Empower Users. *CHI'90*, 337-347.
4. Fischer, G., K. Nakakoji, J. Ostwald, G. Stahl and T. Sumner. Embedding Computer-based Critics in the Contexts of Design. *INTERCHI'93*, 157-164.
5. Goffman, E. Presentation of Self in Everyday Life. Anchor, 1959.