Using Shared Knowledge for Online Photo Access Control

Michael Toomim
University of Washington
toomim@cs.washington.edu

Xianhang Zhang
University of Washington
xianhang@u.washington.edu

ABSTRACT
We propose a novel security mechanism for controlling access to socially-sensitive content, such as online photographs. Rather than requiring entry of explicit access control lists, we allow users to control access implicitly through shared knowledge questions. We conducted a study that describes users' privacy concerns about online photo sharing, and shows users readily understand the concept of shared knowledge questions and can design such questions with a reasonable amount of effort.

ACM Classification: K6.5 [Management of Computing and Information Systems]: Security and Protection. - Authentication.

General terms: Security, Experimentation

Keywords: Privacy, Access Control, Photo Sharing

INTRODUCTION
People are increasingly sharing photos and other personal artifacts online, but one might prefer their boss, family, or a stranger not see some photos. To do this, website users must specify *access control*: a set of rules that allow access from some people, and deny it from others.

Current photo sharing websites use variations of *whitelists* and *blacklists*, where users explicitly list people or groups who should or should not get access. However, as we describe below, white and blacklists are tedious and inflexible, and can be rude.

We propose, instead, that sharers design guard questions such as "where did I travel this summer" or "what is my dog's name" that must be answered to view a photo or album, leveraging the shared knowledge pre-existing in social networks. We conducted a study to investigate the design of guard questions, including the effort it takes to design them, and the types of questions sharers use to guard for different classes of people.

PROBLEMS WITH WHITELISTS AND BLACKLISTS
Whitelists and blacklist require users to explicitly translate social relationships into lists of account names and/or email addresses. This is problematic in a few ways:

Tedious
Creating and maintaining lists for many photos or albums, each with many sharees, requires substantial tedious work, particularly for people without existing website accounts, and makes it easy to forget to include people.

UIST'07, October 7–10, 2007, Newport, Rhode Island, USA.
ACM 978-1-59593-679-2/07/0010.

Inexpressive or complicated
To alleviate the tedium of large lists, websites let users white or blacklist predefined groups of users, such as "close friends" or "family". However, these mechanisms cannot express some access requirements. For example, if I wanted to hold a surprise birthday party for John, the group "All my friends except John" is impossible to express on every major photo sharing site.

On the other hand, more expressive grouping mechanisms, such as those in operating systems, become complicated to use in ways similar to programming: they require education, abstract reasoning, advance planning, and debugging.

Thus, white and blacklists exist in a bounded sea of zero-sum tradeoffs: without groups they are tedious, with arbitrary groups they are complicated, and with predefined groups they are inexpressive. Guard questions may be more flexible.

Rude
Social relations are inherently soft and ambiguous, yet white/blacklists are hard and binary. It can be insulting to learn you are on a friend's blacklist. We suspect it could be less insulting to be unable to answer a question about her summer travels. As a medium, the internet already polarizes social relationships and it is worth pursuing authentication policies that allow more social nuance.

USER STUDY
Our study provides a first look at the usage of guard questions, abstracted from any particular implementation. We recruited a total of 11 people to provide a total of 46 photos that they wanted to share with some people, but not have accessible to everyone. Users reported who they would want and not want to see each photo, as well as the importance of these people on a 4 point Likert scale. Finally, they designed guard questions that they felt would effectively control access to each photo, and described the difficulty of the process.

RESULTS
Who we want to share photos with
We collected 78 responses which we clustered into 5 emergent categories:

Type of person	Portion of Sample	How much you want them to see
Friends	43%	1.8
Family	17%	1.8
Specific People mentioned by name (eg: Sam)	19%	2.8
People in the photo/event	17%	2.3
People who know the people in the photo	8%	1.5

Table 1: A list of people which subjects wanted to see the photos

Who we don't want to share with

We collected 67 responses, and categorized them into 6 emergent categories:

Type of person	Portion of Sample	How much you don't want them to see
Authority Figures	30%	3.4
Family	22%	3.0
Strangers	27%	1.9
Specific People	13%	3.5
Friends	5%	2.3
Potential Boy/Girlfriends	5%	1.7

Table 2: A list of people which subjects didn't want to see the photos

Guard Questions

Subjects easily understood the concept of guard questions. They could readily create guard questions after reading a one paragraph description. Our subjects designed 47 different guard questions in which we found 5 categories:

Question Type	Example Question	Frequency
About themselves	What's my favorite spirit for making mixed drinks?	28%
Knowledge of a mutual friend	What was the name of Susan's hairy dog?	23%
About a specific place or event	In what country did I work in Europe?	17%
About something in the photo	What did we do after we left the bar?	17%
General Knowledge	The "AP" in AP Stats stands for?	11%

Table 3: Types of questions our subjects designed

It took between 35 and 100 seconds for most subjects to design guard questions, according to their self-report. This may be similar to the design-time for whitelists. Subjects also gave the following more detailed feedback:

- *the larger the inclusive group the harder it is to make the question*
- *I took longer to think of questions for pictures that I only wanted specific people to know, because it was harder to think of specific things that only they would know."*
- *Question were generally easy except where picture was of me only*
- *It seemed like I had very few groups that I needed to segregate. Thinking too hard about groups wasn't really worth it.*
- *it took quite a bit of thought -- I had to go back and forth before I was somewhat confident that I had a decent guard question. the questions that were hard were with more open groups of people instead of guarding one person.*

Our results also show that the relative frequency of different guard questions varied depending on who the user was trying to exclude. The correlations, such as excluding family via questions about friends and events, have some intuitive appeal.

	Self	Friend	Place/Event	Photo	General
Authority	+++	-	-	+	--
Family	---	++	++	++	---
Stranger	-	+++	+	---	+++
Specific	---	--	++	-	+++

Table 3: Relative frequency of different guard questions for excluding different groups. + means more frequent, - means less frequent.

CONCLUSIONS & FUTURE WORK

We have presented a new mechanism for implicit access control based on shared knowledge. We have shown that users find the concept intuitive, that they can design questions with a moderate amount of effort, and that they design different classes of guard questions to exclude different classes of people.

Our subjects were told nothing about the implementation of the guard question verification. As a result, many of the questions would be difficult to verify in an automated system. (*E.g.* "What made the party stop being fun?") One area of future study is to see how requiring computer verifiability would affect the difficulty guard question design.

Another important component is to establish the security of guard questions and to see whether a user's perception of security is matched by actual security. If this turns out to be the case, then guard questions could represent an effective alternative to existing photo sharing mechanisms.

ACKNOWLEDGEMENTS

We thank James Fogarty, James Landay, Tom Furness, and Tadayoshi Kohno for their support and advice. The first author was supported by a National Science Foundation fellowship, and the second by a grant from the Department of Homeland Security.