Photo-Based Authentication Using Social Networks

Sarita Yardi[2], Nick Feamster[1], Amy Bruckman[2]
[1]School of Computer Science, [2]School of Interactive Computing
Georgia Institute of Technology
yardi,feamster,asb@cc.gatech.edu

ABSTRACT

We present *Lineup*, a system that uses the social network graph in Facebook and auxiliary information (*e.g.*, "tagged" user photos) to build a photo-based Web site authentication framework. Lineup's underlying mechanism leverages the concept of CAPTCHAs, programs that are designed to distinguish bots from human users. Lineup extends this functionality to help a Web site ascertain a user's identity or membership in a certain group (*e.g.*, an interest group, invitees to a certain event) in order to infer some level of trust. Lineup works by presenting a user with photographs and asking the user to identify subjects in the photo whom a user with the appropriate identity or group membership should know. We present the design and implementation for Lineup, describe a preliminary prototype implementation, and discuss Lineup's security properties, including possible guarantees and threats.

Categories and Subject Descriptors

H.5.m [**Information Interfaces and Presentation**]: Miscellaneous

General Terms

Design, security

Keywords

Social networks, trust

1. INTRODUCTION

Many Web sites need lightweight authentication schemes to distinguish human from non-human users or to control distribution of content to select groups. However, today's Web access control mechanisms remain fairly cumbersome; administrators must maintain access control lists and user accounts, and users must remember and manage a large collection of passwords. An authentication mechanism known as CAPTCHAs has gained popularity for distingishing humans from non-humans. CAPTCHAs present automatically generated graphical images to a user that contains some text and asks the user to identify the sequence of characters that is presented in the graphic. Although CAPTCHAs are helpful for distinguishing humans from non-humans ("bots"), they do not attempt to verify an individual user's identity. CAPTCHAs may be cumbersome in a number of other ways: (1) they may be too obscure for even a human to decipher (i.e., they are not usable); (2) large numbers of CAPTCHAs and their solutions must be generated.

In some cases, Web site administrators need more fine-grained access control over their sites, such as the ability to determine whether the client attempting to access the page is a member of a certain group. To address such cases, we propose Lineup, a lightweight authentication mechanism to help administrators identify a user's identity or membership for site-level access using CAPTCHA-like mechanisms. Lineup relies on a simple idea: present photos to a user and ask that user to identify the names of subjects in those photos. This seemingly simple mechanism presents the following challenges:

- To authenticate a user's identity, the photo presented to the user must include a group of people uniquely known to that user. This requirement also implies the need to identify human subjects in photos (*e.g.*, creating the "solution" to a Lineup photo).
- To authenticate a user as a member of a certain group, Lineup must be able to determine that the subjects in photos are also members of that group.
- Lineup should present tests to the user that the user can answer, which may require presenting the user with a series of tests.

To solve these problems, we propose that Lineup use the Facebook social network as a substrate for building such trust relationships. Facebook is a social networking site where users create profiles and connect to their friends. Individuals' Facebook networks often reflect their real-world social graph more closely than they do in related sites such as MySpace[1] and Orkut[2]. In contrast to these less-structured sites, the technical and social design of Facebook encourages users to articulate existing relationships by joining networks, groups, and filling out profile fields [4]. This articulation of one's real-world networks might help establish some level of

[1]http://www.myspace.com
[2]http://www.orkut.com

WOSN'08, August 18, 2008, Seattle, Washington, USA.

accountability among Facebook users, suggesting that they will be less likely to engage in deceptive practices. As such, we hypothesize that an inferring trust within Facebook could be more effective and accurate than on other social networking sites.

Facebook and its open API provide the necessary elements for building an authentication framework by enabling access to a user's friend links and group membership. More specifically, Facebook users upload photographs and manually "tag" photos with the identities of the subjects in the photographs. Lineup exploits the strong relationships that exist in Facebook's user graph to infer this trust with some level of confidence. The system treats the existence of a link between two people as an indicator that the two people know (and, to a certain extent, trust) one another. The system then makes more fine-grained inferences of trust levels based on user-specifications. To authenticate a user, Lineup presents a photograph containing people purporting to be that user's friend, based on shared links and memberships within their social graph. If the links are in fact accurate representations of real-world friendships, then it is likely that a user who is presented with photographs containing these relationships can identify subjects from these relationships in the photograph.

We have built a prototype implementation of Lineup modeled as a standard client-server architecture. A Web server presents one or more Facebook photos to a user (the client) and requests that she identify people or characteristics of the photos to authenticate herself and gain access to the site. Lineup uses the structural group-level features of Facebook to infer basic levels of authentication. This is a relatively weak form of authentication; we do not envision that it will be used for highly secure environments, such as bank accounts. However, for many popular Web sites, such as blogs, wikis, and photo-sharing portals, these levels of authorization are likely to offer users sufficient granularity and control, while also providing them with a simple mechanism to leverage their existing social graph. This reduces the overhead of having to assert individual user-level access and having to implement access control on a site by site basis.

The rest of the paper is organized as follows. In Section 2, we briefly describe related work on image-based authentication. Section 3 describes the design goals of our system and in Section 4 we describe design specifications. In Section 5 we describe our prototype implementation. In Section 6 we discuss security issues and finally, in Section 7, we address possible challenges and conclude with future directions.

2. RELATED WORK

Recent measurements of online social networks have observed a number of shared structural characteristics of these networks, such as power law distributions, scale-free properties, graph evolution, and information cascades [8, 2, 5]. Although these studies frequently conclude with implications of findings for future work, we note that to date, there are few examples of real-world applications that leverage the explicitly social nature of users and content in a graph for building trust into distributed systems. Some studies have modeled real-world networks, such as blogs and viral outbreaks (*e.g.*, [6]), but these focus on networks as a function of nodes and paths, rather than a function of the subjective and individual relationships and content. In this section, we review structural properties of social networks in the context of Facebook, and suggest their implications for authentication and trust inference in our system.

We base the design of Lineup on the concept of CAPTCHAs, automated tests that take advantage of human processing power to differentiate humans from computers [1]. CAPTCHAs are most commonly used to authenticate users by requiring that the user type the letters of a distorted sequence of letters or digits. CAPTCHAs have also been proposed for other important practical applications, such as image search, content filtering, spam, common-sense reasoning, computer vision, accessibility, and security [1]. Although CAPTCHAs have been incorporated into a number of the above applications to authenticate a visitor as human, rather than machine, CAPTCHAs have rarely been explored as mechanisms for authenticating a user's identity *relative to* other human users. Can CAPTCHAs be used to positively distinguish one person from one another, rather than just a person from a machine? A mechanism with this feature opens new research directions in a number of domains, including assessing trust in Web site authentication based on individual user-access controls.

One reason human-human authentication has not yet gained traction in comparison to human-computer authentication is that the process of distinguishing between one person and another requires specific knowledge and context about each individual. There must exist a predefined set of taggings of any given image from which to test one user's identity in comparison to another, but the process of automating knowledge generation—a "smart" CAPTCHA—of individuals is difficult, if not impossible, with current technologies. Furthermore, manual tagging of photos is laborious and not scalable. Crowdsourcing has been popularized as an approach to overcome this challenge, but its real-world feasibility is not yet well-documented (*e.g.*, see Googe Image Labeler[3]).

3. DESIGN GOALS

We intend Lineup to be used as a lightweight authentication framework to enable content publishers (*i.e.*, users who post content to the Web in the form of blogs, photo albums, etc.) the ability to easily and flexibly control access to their published content. Lineup is not intended as a foolproof mechanism; it prioritizes ease-of-use and tractability over more secure password-based authentication. Any user who can identify the subjects in a photograph, and optionally some additional contextual information, can gain access to the restricted content.

However, Lineup's authentication provides several ancillary benefits that might cause a publisher of content to opt to use Lineup in lieu of standard password-based authentication. In particular, Lineup offers the following benefits:

- *Flexibility for administrators/publishers.* Content publishers may wish to flexibly and dynamically update the group of users who have access to a set of content.
- *Convenience for clients.* Users may not always want to remember passwords or accounts for access to every new site or service, particularly when a publisher may want to make each collection of content (*e.g.*, photos for a particular party, a blog entry aimed at a particular audience) accessible to a different group of users,

[3]http://images.google.com/imagelabeler/

and when the membership in the groups themselves may be dynamic.

- *Large number of challenges with readily-available solutions.* Asking clients to solve puzzles such as CAPTCHAs requires that the server have at its disposal a large number of challenges (in our case, photos) with readily available answers. Facebook photos can serve as a ripe source for such tests, due to the large volume of tagged photographs that are uploaded every day.

4. SYSTEM DESIGN

The social tagging process in Facebook introduces a new large-scale phenomenon that we can exploit to automate authentication with a reasonably high level of confidence. Facebook is the number one photo sharing application on the web, with over 14 million Facebook photos uploaded daily (more than the next three largest photo-sharing sites combined)[4]. Furthermore, there are over 66 million users and 250,000 new users a day, with more than 2.2 billion tags in these photos[5]. Although not all users upload photos, and not all of that subset tags photos, the sheer size and scale of overall activity suggests that the user base may be large enough to authenticate a large number of users.

Lineup is implemented as a Facebook application widget that can be embedded in 3rd-party Web sites, such as LiveJournal, MediaWiki, or Flickr, using a simple "embed" script. To use Lineup, a user simply fills in application parameters, including the site that she wishes to control access to, and the Facebook group, network, or event(s) to whom she wishes to give access. Many blog systems like Wordpress or Movable Type support built-in widget management systems as plug-ins. Lineup functions similar to a widget engine, returning an embeddable section of code that she can cut and paste into her site, similar to Google's SearchBar Widget or Yahoo's! Weather Widget[6]. Below we describe Lineup's high-level design and usage model. We describe first how a Web site, or other type of content-providing account, might embed Lineup's authentication mechanism and then describe how the authentication process might incorporate various levels of trust and privacy.

4.1 Design Overview

We intend Lineup to be used for lightweight authentication for a growing number of "semi-public" sites and services. Many potential cases may arise where a user wants to easily restrict access to some content, such as a blog or photo album, to a specific group of users. In these cases, it is not catastrophic if the access control is breached; the primary design goal is to ensure that a user can easily generate access control restrictions based on such groups.

This authentication process proceeds in the following steps:

1. A user defines access control for a certain Facebook group, network, or event. She then associates this access control with some third-party site, such as a LiveJournal blog (*e.g.*, only members or this group, network, or event can have access to this content).

2. The user specifies high, medium, or low privacy settings for the protected site or content.

3. When a client wishes to log in to that site, the site fetches one or more pictures from Facebook that contain members of the associated group that is allowed to access the content. Facebook then returns to the site both the picture and a list of names of people in that group (*i.e.*, the "solution").

4. The user identifies names of users, as well as optional advanced tags, such as location, to gain access to the restricted content.

4.2 Levels of Trust and Privacy

Lineup infers levels of trust that should be measured based on a user or site administrator's desired privacy setting using one or more photos that are in her specified group. We draw from Lessig's regulatory architectures of control that govern online behavior [7]. We propose two basic categories of association for Lineup: *memberships* (*e.g.*, groups, networks) and *events*. Photos within both categories can contain one or more types of tags, which can then be used as authentication mechanisms. Additionally, for a particular *event* type, Lineup can ask a client: (1) Who was there? (2) When was it?; or (3) What happened? For a *membership* type, Lineup can ask: (1) Who is a member? (2) What is the name of it? (3) Why was it created?

In both cases, the first requires minimal knowledge of the environment and fulfills only a low privacy requirement. The second requires slightly more familiarity, although could still be guessed by a close outsider. The third is the highest privacy setting and is set by the owner of the third party site itself. She might require that the user describe an obscure object in the photo, an event that was occurring during that photo, or other people who were nearby when the photo was taken. Thus, although not ideal, the burden is on the application user to establish how high the highest privacy setting is. Similar authentication mechanisms have emerged in some smaller, niche-based sites. For example, Sconex[7] is a social networking site for high school students that requires that students answer questions about the school with which they claim affiliation: "Who teaches 10th grade English?" and "What color are the first floor lockers?" are simple for students to answer but more difficult for outsiders.

Users already tag photos with individual identities in Facebook, and Lineup automatically incorporates the second privacy setting level, such as the group, network, or event name, which can be extracted using the Facebook API. Only the third level requires additional effort from the part of the user using the API. Lineup computes a score based on the server-level privacy setting (high/medium/low) and how many tags are accurately identified by the user. A low privacy setting and a correct identification of who is in the photo will authenticate the user. A high privacy setting and a correct identification of who is in the photo but incorrect identifications of when/what/why will reject the request.

4.3 Installing the Authentication Framework

A Facebook user who wishes to implement privacy control settings on an external third party site using our Facebook API first logs in to the application and indicates which site

[4]http://www.facebook.com/press/info.php?statistics
[5]As of April 2007.
[6]http://www.googlewidget.com/googlewidget.html, http://widgets.yahoo.com/widgets/yahoo-weather
[7]http://www.sconex.com/

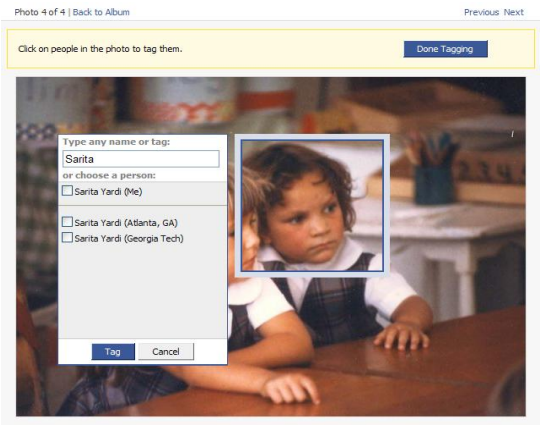

Figure 1: Photo tagging process

```
//User identification
$facebook = new Facebook($appapikey, $appsecret);
$user id = $facebook->require login();

//Get user photos
$photosOfMe = $facebook->api client->photos get($user id);
$getTags = $facebook->api client->photos getTags($photo['pid']);
```

Figure 2: Identify user and photo tag details

```
//Extract photo details
foreach ($getTags as $tag) {
    $imgSrc = $tag['pid'];
    $subject = $tag['subject'];
    $description = $tag['text'];
    $XCoordinate = $tag['xcoord'];
    $YCoordinate .= $tag['ycoord'];
}
```

Figure 3: Extract photo tagging details

she wants to protect. She then has the option of setting high, medium, and low access levels. Trust inference is an an imperfect process and these are not presented as universal privacy settings, but are relative to the site being protected and the Facebook network that she chooses to grant access to. For example, a Flickr user might choose to set high privacy settings to her personal Flickr page, which she has chosen to share with only family members in Flickr. In contrast, she might choose to set low settings to the rest of her Flickr site that she has made openly available to her entire Flickr contact network. Analogously, in Facebook, she might mentally map the high level setting might be to a Facebook group of her closest college friends, of whom there are only five members and who she considers family. In contrast, she might set her low privacy settings to her entire university network, with whom she has only a loose affiliation.

5. PROTOTYPE IMPLEMENTATION

Lineup is implemented using Facebook's PHP library and a MySQL backend database. Each file contains three sections: authentication, database queries, and formatted output. Each page first authenticates itself using the key and password to connect to Facebook, then queries the database using the functions provided in the Facebook library or by directly querying the database using fql, Facebook's modified version of SQL. The display is output using Facebook's built-in FBML language.

5.1 Implementation Overview

To use Lineup, users first upload their photos into Facebook. They are then able to tag the photos with one or more user names (see Figure 1). These are usually names of other users in Facebook, but the system allows any descriptor to be used. Thus, a photo might be tagged with event or group descriptions or other unique identifiers. The system also allows any other user to tag a photograph, but the tagging must be approved by the photo uploader before it is saved. A user then enters her site information into Lineup and sets the access level appropriately. The system first identifies the user using the built in Facebook login function (see Figure 2). The system then pulls in photos from that particular user's photo database. Depending on the access levels set, the photos might be relatively simple photos with only one tag, or more complex and difficult to identify photos. The system can select a particular photo or set of photos using an album id or a photo id, or "pid" (see Figure **??**). Finally, Lineup retrieves the tags of each photo and the coordinates of each tag (see Figure 2). Tag subject, description, and coordinates are compared to client input in the system and either authenticated or rejected based on a match (see Figure 4). Similar to CAPTCHAs, if the user is unfamiliar with a particular event, but feels she should be granted access to the site, she can choose to try another photo. However, too many restarts will lock her out of Lineup. A limitation of Lineup is that it only works for registered Facebook users. Similarly, for photos that contain multiple people, some of whom may be not tagged, the user must additionally specify if all persons in the photo need to be identified or just a particular subset.

6. SECURITY ANALYSIS

A user might be able to determine the identity of users in photos by crawling the Facebook site. Although the ability to crawl the site is limited, the site still typically allows users to see photos of friends-of-friends. A resourceful attacker could perhaps harvest these photos to learn the identities of a group without actually knowing anything about the group or its members. To defend against this attack in general, Lineup would need to apply the third level of privacy (*e.g.*, what is happening in this photograph?). However, in many cases, a user may not be able to cull enough pictures of friends from these relationships to identify all members of a group. Furthermore, such a process requires significant effort on the part of the user, and the effort required would likely outweigh the benefit, given that we only intend Lineup to be used in cases where the "stakes" are low (*e.g.*, photo sharing Web sites, blogs, etc., as opposed to online banking).

Because many users expose a photo of themselves even if their profile is private, a user might also be able to browse Facebook to cull user photos that might help them identify members of a certain group. Similarly, such an attack might require substantial enough effort to thwart most users from attempting to see whatever content is protected.

6.1 Fraud and Denial of Service

A malicious user could also mount various attacks to prevent other users from gaining access to a service. For example, someone who wanted to deny users access to a par-

Figure 4: Photo identification process

ticular group could intentionally mislabel photographs, so that a member wanting to gain access to the group would be presented with photographs containing users that were unknown to the client. Such an attack might be used to deny a client access to a Web site or service. Although more thorough analysis is needed, we posit that this attack would be quite difficult to mount in practice: An attacker would have to both upload and tag a large number of photos to increase the likelihood that a client would be presented with photos from unknown users. Additionally, to preclude such an attack, Facebook could confirm a tagging operation with the individual being tagged in a photo.

6.2 The Imperfect Nature of Friendships

Much of social network research focuses on measuring flow by locating actors in the network—who are the connectors, mavens, leaders, bridges, and isolates? A naïve approach might attempt to model trust based on the existence of individual links between users; however, the semantic meaning of "friend" is highly subjective and contextual and requires a more complex representation. Trust between individuals is contextual, based on shared history and prior interactions, yet is reduced to a relatively unsophisticated binary representation in a social graph. Sociologists have emphasized the importance of establishing trust in online reputation systems that collect, distribute, and aggregate feedback about participants' past behavior [9]. Facebook mimics such a reputation system, in which past shared history motivates future trust. Using the API, we can capture shared activity and interactions between users to assert some reasonable measurement of trust.

Network outliers, however, exist at the individual level. Could the authentication application somehow be hijacked by such outliers to harvest information about who is in photos? For example, a socially excluded individual might try to hack into a site that her peers have access to. She will bypass low level security, being able to identify peers in photos, even if she is not a member of their Facebook group. She may also be able to identify when or where an event took place, or what the name of a group is. Thus, the onus is on the original user who set the privacy settings to set the highest level access tag as something that this outlier cannot guess. Relying on human action for highest level privacy is not ideal and is an ongoing challenge in our design. We are considering ways of employing existing contamination detection methods in graphs to detect the infiltration of social outliers (*e.g.*, [3]).

7. CONCLUSION AND FUTURE WORK

This paper has described the motivation and design of Lineup, a framework for authenticating members of groups using photographs. Lineupis inspired by CAPTCHAs, but provides additional functionality. In addition to simply distinguishing humans from non-humans, Lineup can help a site administrator or content publisher verify a client's membership in a certain group, based on whether the client can identify the subjects in a group of photos. Lineup provides a lightweight authentication mechanism that allows content publishers or Web site administrators to control access to various content using a flexible, easy-to-use graphical authentication framework. Lineup allows publishers to restrict access to content to specific social network groups (*e.g.*, Facebook groups, attendees of an event) without being required to specify individual site-specific access control. Similarly, clients can use Lineup to quickly authenticate to a Web site without having to enter or remember passwords.

Despite its apparent simplicity, many questions remain for implementing Lineup in a real-world social network. First, we need to analyze system performance to determine whether a small group of photos could be used to reasonably distinguish one group of users from another (*e.g.*, it may be likely that the photos for one event or group might overlap highly with those from another gorup, and finding "challenges" that distinguish one group from another may be difficult). We also need to assess usability by testing Lineup with Facebook users. Finally, we must perform a rigorous analysis of Lineup's security properties. We are studying these questions as part of our ongoing work towards developing a real-world deployment of Lineup.

8. REFERENCES

[1] L. V. Ahn. *Human computation*. PhD thesis, Carnegie Mellon University, 2005.

[2] A. Barabasi. Scale-free networks. *Scientific American*, 288:60–69, 2003.

[3] K. Bharat and M. R. Henzinger. Improved algorithms for topic distillation in a hyperlinked environment. In *ACM SIGIR*, pages 104–111, Melbourne, Australia, 1998.

[4] d. boyd. *Why Youth (Heart) Social Network Sites: The Role of Networked Publics in Teenage Social Life*. MIT Press, 2007.

[5] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Trans. Web*, 1(1):5, 2007.

[6] J. Leskovec and C. Faloutsos. Scalable modeling of real graphs using kronecker multiplication. In *Proc. of the 24th intl conf on Machine learning*, Corvalis, Oregon, 2007. ACM.

[7] L. Lessig. *Code Version 2.0*. Basic Books, New York, 2006.

[8] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proc. of the 7th ACM SIGCOMM conference on Internet measurement*, San Diego, CA, USA, 2007. ACM.

[9] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.