

Context-based Personalised Settings for Mobile Location Sharing

Fehmi Ben Abdesslem
and Tristan Henderson
School of Computer Science
University of St Andrews
St Andrews, Fife, United Kingdom
{fba,tnhh}@st-andrews.ac.uk

Sacha Brostoff
and M. Angela Sasse
Department of Computer Science
University College London
London, United Kingdom
{s.brostoff,a.sasse}@cs.ucl.ac.uk

ABSTRACT

Location-Based Services (LBSes) are increasing in popularity, but create many privacy concerns for users. LBSes usually rely on the same default privacy settings for all users. In this position paper, we claim that such settings are inappropriate for location sharing and that settings should instead rely on contextual information to recommend personalised privacy settings for users. We present results of an initial user study ($n=80$) to corroborate this position, and suggest avenues for further research.

1. INTRODUCTION

Location-Based Services (LBSes) have become increasingly popular and, coupled with recent growth in smartphone usage, provide an interesting platform for the deployment of recommendation and personalisation technologies. These services, however, introduce many potential privacy problems: a recent survey reveals that 84% of users are concerned about potential losses of privacy when using LBSes, and 49% would be more comfortable if they could easily and clearly manage who sees their location information.¹ Managing who can see one's personal information is currently done through static settings for most popular Social Network Sites (SNSes) such as Facebook, Twitter or Google+; users are asked to adjust their settings and these are typically applied to all interactions with a given member or set of members on the site. Such static settings might be appropriate for sharing static information: a user's birthday does not change very often, and so the set of users with whom a user might wish to share their birthday might also not change. But this might not be true when sharing varying information such as location, as privacy preferences may depend on the particular context in which a location is visited, and not only on the person with whom it would be shared.

Nonetheless, LBS settings are usually static, not only for Facebook or Google+'s location-sharing features, but also for SNSes dedicated to location sharing such as Gowalla or Foursquare. This may be due to the disruption that would be

caused by asking a user about their sharing choices for every location. Alternatives to such a disruptive solution would be to automatically predict a user's sharing choices, or to recommend privacy settings to the user.

In this position paper, we propose that before we attempt to provide personalisation and recommendation technologies to mobile users, we should first develop personalised privacy settings that can better cope with users' actual sharing preferences than static settings. Using data collected from a mobile location-sharing study with 80 participants, we investigate whether the personalisation of settings can leverage contextual and self-reported information.

The remainder of this paper is organised as follows. In the next section, we demonstrate how static settings are not appropriate. Section 3 defines what context could be exploited by LBS settings and how it could help predicting users' location sharing choices. We then review related work in Section 4 and discuss our position in Section 5.

2. STATIC PRIVACY SETTINGS

Static settings for LBSes would be appropriate if users actually had static privacy preferences while sharing their locations with other people. To determine just how constant users' behaviours are, we conducted an experiment with 80 student participants, each carrying a mobile phone for seven days and sharing their locations with their social network.

Our methodology [2] is divided into three consecutive phases:

- **Pre-briefing.** Participants install a Facebook application on their personal Facebook account and fill in a questionnaire.
- **LBS usage.** Over the course of seven days, participants use *LocShare*, a custom mobile application allowing them to share their location with their social network on Facebook. Using the Experience Sampling Method, they also reply to automatically generated questions about their experience and sharing choices, sent to them on the mobile phone.
- **Debriefing.** Participants were interviewed and asked for more explanation about their sharing choices.

¹Microsoft Research Data Privacy Day 2011, <http://www.microsoft.com/privacy/dpd/>

Participants were periodically asked to share their current location, or a picture of their current location. Upon receiving a sharing request, a participant had the choice of replying or ignoring the request.² In addition to these prompted events, *LocShare* allowed participants to share their location (described with text comments or a photo) of their own accord without being prompted.

In total, 4,334 sharing requests were sent to participants, which led to 2,064 replies (a 47.6% response rate). Participants also shared 531 locations without prompting. We refer to any of these 2,595 events as a *sharing event*. At each sharing event, participants had the choice of sharing their location with: everyone on Facebook, only their Facebook friends, only a subset of their Facebook friends, or no-one. In the latter case, the location was still uploaded on Facebook but remained only visible to the participant.³ We further refer to any of these choices as a *sharing choice*.

Looking at each of the 80 participants, we see that none of them made the same sharing choice at each of their sharing events. In other words, not a single participant would have been satisfied with static privacy settings for their location-sharing application. To verify this assumption, we asked participants in the pre-briefing to choose a default sharing choice for the experiment, that would be used when the system decided to share their location without prompting. Figure 1 shows the Empirical Cumulative Distribution Function of the sharing choices made by the participants. 23 participants (28.8%) never used their default sharing choice for their sharing events. The participant who most used the default sharing choice used it for 97.6% of their sharing events. Hence, all of the participants made at least one sharing choice that differed from their static default sharing choice. From Figure 1 we also see that the median participant used their default settings for just 9.5% of sharing events.

In the pre-briefing, participants were asked whether they had changed their Facebook privacy settings from their default values. Interestingly, Figure 1 shows that those participants who use the default Facebook privacy settings (23 participants) were more likely to use their default settings during the experiment: the median participant in this group used their default settings for 33.3% of their sharing events. On the other hand, the median participant from the group that used personalised Facebook privacy settings used their default settings for only 7% of their sharing events.

Not only are static privacy settings inappropriate, they may lead to unintended information disclosure. Of the 1,889 sharing events where participants did not use their default sharing settings, 455 of the events (24.1%) used settings less

²Ignoring a request was not always deliberate. Participants may have not heard the mobile phone ringing, or were temporarily not covered by the cellular network. The sharing request would then timeout and no longer be displayed on the phone.

³To test a different hypothesis [8], half the participants were allocated to a simulated sharing group, where locations were not actually shared, but participants were asked to act as if they were. Results from the two groups are merged in this paper.




Figure 1: Empirical CDF of the proportion of sharing events where a participant shared their location using their default settings. No participant made the same sharing choice at every sharing event.

permissive than the defaults. In other words, if participants were unable to personalise their default privacy settings, using their static default privacy settings would have incurred 455 privacy leaks, by sharing their location to people with whom they were not willing to share. We therefore believe that *LBSes should provide personalisable privacy settings that reflect users’ personal needs*.

3. TOWARDS CONTEXT-BASED SETTINGS

We have seen that location privacy settings are not static in practice. We believe that this is because they depend on a user’s current context, e.g., where the user is, who they are with, what is nearby, and so forth [9]. Indeed, personalised *LBSes* are a good example of the context-aware application as defined by Schilit et al. [9], as they may involve *Automatic Contextual Reconfiguration*, as the *LBS* is reconfigured according to the context, and *Context-Triggered Actions*, as the context might trigger an action from the *LBS* (e.g., recommending new privacy settings for users).

We can imagine a multitude of simple examples where context might influence Alice’s decision to share her location. In a first scenario, Alice does not mind sharing her location to her friends and family wherever she goes, but some evenings, she is tired and does not want them to pop-up when she is at home. She would then hide her location for the evening. In another scenario, Alice does not want her family to know that she is with Bob at night. Since she has no control over Bob’s sharing choices, she wants to hide her location to her family in the evening when she is with Bob, as otherwise Alice and Bob would both share the same location and therefore disclose that they are seeing each other.

While Alice’s desired privacy settings are not static, asking for her location-sharing preferences at every new loca-




Figure 2: Private and public rates for different types of locations. Participants were less likely to share their location in Library and Residential areas.

tion would be highly disruptive. An alternative solution is then for the LBS to infer Alice’s privacy settings according to the current context. To do so, the system needs to sense context, and then determine which contextual information might influence Alice’s sharing choices.

Recent research has shown that automatically sensing context is possible, for instance by sensing activity using the accelerometer and microphone [6], location using GPS/Wi-Fi/GSM but also compasses and accelerometers [3], and co-located people using Bluetooth [7]. Hence we believe that context can be used by LBSes. But using this sensed context to determine what information influences users’ sharing choices introduces new challenges. The particular information may vary for each user and may be based on complex social considerations.

So if context can be used to recommend and personalise privacy settings, what contextual information should be used? We now investigate the contextual and self-reported information that influence our participants’ location sharing choices. We study two extreme behaviours: sharing with nobody, and sharing with everyone on Facebook. We define the *private rate* as the proportion of events where locations were shared to no-one, and the *public rate* as the proportion of events where locations were shared to everyone.

Location itself is the first obvious piece of contextual information that may influence user behaviour. Figure 2 shows that both private and public rates depend on the type of location. Participants have higher private rates in the Library and Residential areas, and thus share their location less than in other types of places. On the other hand, public rates are higher in Leisure, Academic and Retail places. During the debriefing interviews, some of the reasons provided for not sharing were that participants did not want people to join when they were at the Library, and did not want them to know that they were staying home on a Saturday night.




Figure 3: Private and public rates when the participants are alone, with friends, and with strangers. Participants were less likely to share their location when alone.




Figure 4: Private and public rates during day time (between 0800 and 2000) and night time. There was little difference in sharing behaviour over time.

The presence of other co-located people is another example of contextual information that can influence location sharing. At each sharing event, participants were asked if they were with no-one, with friends, or with strangers. Figure 3 shows that the private rate is higher when the participants are alone than when they are with strangers.

That said, not every aspect of the context appears to influence sharing choices. For example, Figure 4 shows no clear difference between day time (0800-2000) and night time.

Nonetheless, we believe that these initial results indicate that gathering more empirical data about users’ choices and corresponding contexts can provide a better insight on what contextual information could be used to recommend personalised privacy settings to LBS users.

4. RELATED WORK

Context is key in the development of new services that will impact social inclusion for the emerging information society [4]. It has already been used for various purposes, such as predicting users' interests for websites [11]. Anthony et al. [1] show experimentally that privacy preferences vary with place and social context. Fang et al. [5] exploits this observation by proposing a privacy recommendation wizard for Facebook settings. By asking limited questions to users about what they would disclose to some of their friends, they generate personalised privacy settings. Toch et al. [10] use the same concept for Locaccino, an LBS that allows the users to set their default privacy settings by defining who, where, and when they want their location to be disclosed.

5. WHAT NEXT?

The rise of location-based services will enable the application of personalisation and recommendation technologies in people's everyday lives. But if people are to use these services, their requirements need to be respected. Static default privacy settings provided by current location-based services are inappropriate as they do not meet these requirements, where the privacy settings for a particular event may be determined by the context in which the event takes place. We have outlined these user requirements by collecting data through a user study, and conducted some initial analysis to show how we might build LBS systems that can recommend personalised privacy settings.

Building a usable LBS with personalised privacy settings presents many interesting challenges. One approach is to bootstrap the system by asking the users for their privacy settings when a new context is detected, and then use these same settings when the same context is later encountered. This solution is highly disruptive at the beginning, but should eventually become less disruptive when no new contexts are encountered. A second approach is to use data collected from other users with a similar profile who have experienced similar contexts. This requires collecting a large amount of empirical data, however, and either determining how to match similar profiles and contexts, or how to choose one default per context. Self-reported data are also a useful source of information about the users' behaviour. For instance, we observed that the participants who reported using Facebook default privacy settings were using default settings more often in our experiment than participants who reported using custom privacy settings on Facebook.

Using a combination of all these approaches, a system could start applying a user's chosen default setting, and then use machine learning techniques to recommend personalised settings, based on the context sensed by the mobile device, on behavioural data collected from other users, and on self-reported information about the user.

Thus, in this paper we argue that before we can apply personalisation to the services themselves, we should apply personalisation to the privacy settings for these services. Such

personalised settings are not only necessary for meeting user requirements, but they should also be considered as a prerequisite for offering personalised services to LBS users: personalised services should not exploit locations that users do not want to share.

6. ACKNOWLEDGEMENTS

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/G002606/1].

7. REFERENCES

- [1] D. Anthony, T. Henderson, and D. Kotz. Privacy in Location-Aware computing environments. *IEEE Pervasive Computing*, 6(4):64–72, Oct. 2007.
- [2] F. Ben Abdesslem, I. Parris, and T. Henderson. Mobile experience sampling: Reaching the parts of Facebook other methods cannot reach. In *Proc. Privacy and Usability Methods Pow-Wow (PUMP)*, Sept. 2010.
- [3] I. Constandache, R. Choudhury, and I. Rhee. Towards mobile phone localization without war-driving. In *Proc. IEEE INFOCOM*, Mar. 2010.
- [4] J. Coutaz, J. L. Crowley, S. Dobson, and D. Garlan. Context is key. *Communications of the ACM*, 48(3):49–53, Mar. 2005.
- [5] L. Fang, H. Kim, K. LeFevre, and A. Tami. A privacy recommendation wizard for users of social networking sites. In *Proc. CCS '10*.
- [6] G. B. Gil, A. Berlanga, and J. M. Molina. Physical actions architecture: Context-aware activity recognition in mobile devices. In *Proc. Workshop on User-Centric Technologies and Applications*, pages 19–27, Apr. 2011.
- [7] V. Kostakos, E. O'Neill, A. Penn, G. Roussos, and D. Papadongonas. Brief encounters: Sensing, modeling and visualizing urban mobility and copresence networks. *ACM Transactions on Computer-Human Interaction*, 17(1):1–38, Mar. 2010.
- [8] I. Parris, F. Ben Abdesslem, and T. Henderson. Facebook or Fakebook?: The effect of simulation on conferencelocation privacy user studies. In *Proc. Privacy and Usability Methods Pow-Wow (PUMP)*, Sept. 2010.
- [9] B. Schilit, N. Adams, and R. Want. Context-Aware computing applications. In *First Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 85–90, Dec. 1994.
- [10] E. Toch, J. Cranshaw, P. H. Drielsma, J. Y. Tsai, P. G. Kelley, J. Springfield, L. Cranor, J. Hong, and N. Sadeh. Empirical models of privacy in conferencelocation sharing. In *Proc. Ubicomp 2010*, pages 129–138, Sept. 2010.
- [11] R. W. White, P. Bailey, and L. Chen. Predicting user interests from contextual information. In *Proc. SIGIR '09*, pages 363–370, 2009.