When Are Users Comfortable Sharing Locations with Advertisers?

Patrick Gage Kelley, Michael Benisch,
Lorrie Faith Cranor, Norman Sadeh

October 2010

CMU-ISR-10-126
CMU-CyLab-10-017

Institute for Software Research
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213-3890

As smartphones and other mobile computing devices have increased in ubiquity, advertisers have begun to realize a more effective way of targeting users and a promising area for revenue growth: location-based advertising. This trend brings to bear new questions about whether or not users will adopt products involving this potentially invasive form of advertising and what sorts of protections should be given to users. Our real-world user study of 27 participants echoes earlier findings that users have significant privacy concerns regarding sharing their locations with advertisers. However, we examine these concerns in more detail and find that they are complex (e.g., relating to not only the quantity of ads, but the locations they receive them at). With advanced privacy settings users stated they would feel more comfortable and *share more information* than with a simple opt-in/opt-out mechanism.

This work has been supported by a Siebel Scholarship, an NSF IGERT grant, and NSF grants CNS-0627513, CNS-0905562. Additional support has been provided by Nokia, France Telecom, Google, the CMU/Microsoft Center for Computational Thinking, ARO research grant DAAD19-02-1-0389 to Carnegie Mellon University's CyLab, and the CMU/Portugal Information and Communication Technologies Institute.

Keywords: privacy, user studies, advertising, expressiveness, location-based sharing

Contents

1 Introduction

In February 2010, the United States House of Representatives Subcommittee on Commerce, Trade, and Consumer Protection held a joint hearing on location information for commercial purposes. In chairman Bobby L. Rush's opening statement, he said:

> "To some extent, location-based services can be viewed as a sub-category of behavioral tracking in that they can quickly, and cheaply, tell advertisers more than contextual advertising ever could about someone's preferences, habits, and patterns." [11]

With intelligent mobile devices growing in popularity, it is possible for future ads to commonly be targeted based on one's location. Google's and Apple's recent purchases of AdMob and Quattro Wireless, are examples of two major technology companies pushing into mobile and location-based advertising.

Additionally, as location-sharing continues to grow as a social phenomenon (e.g., the recent launch of Facebook's Places continues to move this trend towards the mainstream), we have already started to see location-based coupons being offered to users. Foursquare, a popular mobile location-sharing application, is currently leading this push by allowing small businesses and national chains (e.g., Starbucks) to offer recurring, frequency-based, and loyalty-based coupons to over three million users [9]. Businesses that participate in this program are also given access to certain personally identifiable statistics, such as the most recent and frequent visitors to their venues, without users' awareness.

As this commercial landscape expands, we ask what types of privacy settings should be given to users to encourage broader adoption? While location-sharing with friends, family members, and members of social networks has been studied in great detail by researchers [1, 3, 8, 10, 12], the controlled sharing of location information with advertisers is largely unstudied.

In this paper, we present an empirical investigation of peoples' attitudes towards location-sharing with mobile advertisers. By analyzing three weeks of location audits, combined with survey responses, we show users' strong privacy concerns may hinder the adoption of systems leveraging this potentially invasive form of advertising. However, we also find that advanced privacy settings may help alleviate some of these concerns.

2 Related Work

Many research groups have developed location-based services, including PARC's Active Badges [17], Barkhuus's Active-Campus [1], Intel's PlaceLab [5], Carnegie Mellon's Locaccino [12], and Burghardt's Mobile PET work [3]. While much of this work was hampered with adoption and technical issues in its infancy, more recent work has successfully investigated deployed systems where users' locations are queried in the real world. These systems, and their commercial counterparts, have almost always included some sort of privacy module, since even in early wireless advertising efforts it had been reported that participants had privacy concerns regarding location-based advertising [4]. Yet, while this module was always present in discussions [7], it has normally been regarded as a simple "opt-in" or "opt-out" mechanism. Even in focused consumer inquiries, the discussion is largely focused on whether or not users would be willing to opt-in [15, 16].

More recently, some systems have begun to provide more privacy protective measures than simple opt-in mechanisms, including: per-person settings, granularity/resolution control, invisibility modes, and even time and location based rules. A survey of 89 commercial location-sharing applications showed that, in 2009, more than half had some form of privacy control, with varying complexity [14]. However, in all of these commercial and research applications the privacy controls protect users from their friends, not the companies themselves or advertising partners.The bridge between applying advanced privacy controls for location-sharing with friends, family members, and social networks has not been considered for location-sharing with advertisers. Our work unifies these two bodies of work, investigating when consumers and advertisers would benefit from these advanced models.

3 Methodology

In November 2009, we conducted a study requiring users to carry a Nokia N95 smartphone equipped with our location-tracking software, and use it as their primary phone for three weeks. Once a day, users were required to visit our web application where they were asked whether or not they would have been comfortable sharing each location they visited with a number of pre-defined groups, one of which was advertisers. For a complete description of the methodology and technology used in executing this user study (including screen shots of the web application) please refer to [2].

Users were recruited with fliers and newsgroup postings within our university community. Our users were required to have an AT&T or T-Mobile phone contract, and an unlimited data plan (or the ability to obtain one). We compensated participants $30 for their participation and an additional $20-$30 to reimburse their data plan expenses. While we had 111 people fill out our pre-study survey, only 27 users completed the study (due to ineligibility and drop-outs).

Each time a user visited our web application, it iterated through his or her recent locations in order, and it displayed each location on a map. For each location it then asked, “Would you have been comfortable sharing your location **during this time span** with Advertisers?” (where “during this time span” was replaced with the actual time the user visited the location). Users could then respond “Yes, during this entire time,” “No, not during any of this time,” or “Yes, during part of this time. . . .” Users completed these questions for each location at which they spent more than 15 minutes. If a user visited a location at two different times on the same day, it was audited separately for each visit.

4 Results

Our results are broken into two sections: a summary of user-reported information from the pre- and post-study surveys, and a quantitative analysis of users’ daily audits concerning sharing their locations with advertisers.

4.1 Survey Results

Our study included 27 participants (20 male, 7 female), all affiliated with our university community (15 undergraduates, 10 graduate students, 2 staff members). While our users tended to be young (average age was 22.9), more tech-savvy, and most reported daily use of Facebook, we believe this community is representative of early adopters of the technologies that are likely to expose them to location-based advertising.

4.1.1 General Privacy Concerns

Our post-study survey sought to capture users' attitudes relating to two topics: general privacy concern and concerns related specifically to advertising. A summary of user responses related to general privacy concerns can be found in Figure 1. Overall, users report strong concerns across the six questions we asked. The strongest response was in regards to corporate information exchange, we stated: "Online companies should never share personal information with other companies unless it has been authorized by the individuals who provided the information," and 19 of 24 users responded with a 7 on a 7-point Likert scale, representing a strong agreement.

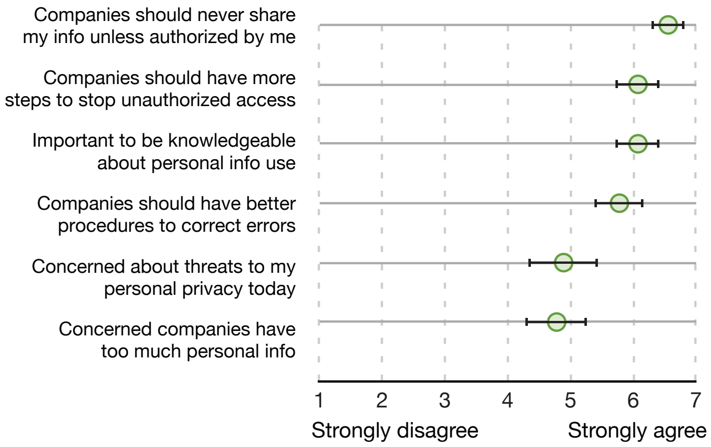

Figure 1: User responses about privacy concerns after the three week study. Answers were reported on a 7-point Likert scale, from 1 (Strongly disagree) to 7 (Strongly agree). Averages and 95% confidence intervals are shown.

4.1.2 Privacy Concerns Regarding Advertisers

The pre-study survey contained one question related to advertisers, asking participants to "rate how comfortable you would be if **advertisers** (e.g., in order to send you promotions or coupons) could view your location," either always, at user-specified times, or user-specified locations. On a 7-point Likert scale, where 1 was labeled "Not comfortable" and 7 was "Fully comfortable," users reported an average of 2.6 for always, 3.6 with specified times, and 4.3

with specified locations. Both time and location specifications made users significantly more comfortable ($p < 0.01$ for both, paired t-tests, time: $t = 3.11$, $df = 24$; location: $t = 4.28$, $df = 24$) and location specifications were significantly more comforting than time ($p < 0.01$, $t = 2.98$, $df = 24$).

The post-study survey contained questions related to advertisers. The first asked "how bad" it would be if a user's location was disclosed to advertisers when they did not want it to be, and also the reverse (i.e., a non-disclosure when disclosure was wanted). On a 7-point Likert scale, where 1 was labeled "Not bad at all" and 7 was "Very, very bad," participants reported an average discomfort level of 1.67 for mistakenly not-disclosing, and an average discomfort of 4.74 for mistakenly disclosing a location. These results suggest that, as expected, a missed opportunity is only a minor concern to our users, whereas disclosing a privacy-sensitive location to advertisers has a significantly higher cost.

We also asked users what the most important factors would be in allowing advertisers access to their locations. The results from this question, again on a 7-point Likert scale from "Not important" to "Very, very important," are displayed in Figure 2. From the reported responses, a user's location and the quantity of ads received mattered significantly more than the brand of the advertisers and time of day.

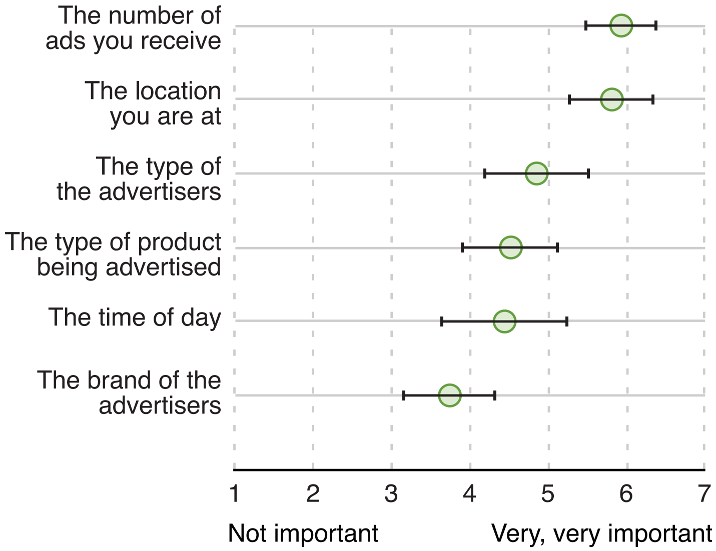

Figure 2: User responses on qualities of advertisers which would impact their future location-sharing decisions. Answers were reported on a 7-point Likert scale, from 1 (Not important) to 7 (Very, very important). Averages and 95% confidence intervals are shown.

4.2 Quantitative Results Based On Location Audits

With over 7,500 hours of audited locations, we sought to quantify when users would be willing to share their real locations with advertisers, and how complex their preferences are regarding this sharing.

In general, users were significantly more willing to share their locations on weekdays from 9:00am–5:00pm (average of 47.5% time shared during business hours compared to 35.8% at other times). The subjects were also significantly more willing to share their second and third most visited locations than their most visited location, which were likely their homes (average of 29% time shared at first most visited, compared to 55% and 41% at second and third most visited locations, respectively).

For the remainder of the analysis, we will compare the following six privacy-setting mechanisms based on the percentage of time our users would have shared under each:

- *Opt-in* – the most common privacy-setting. It simply allows a user to opt-in or opt-out of sharing with advertisers.
- *Time* – slightly more complex than Opt-in. A single rule specifies a time span during which a user is willing to share his or her location (e.g., between 8am-5pm).
- *Time with weekends (Time+)* – a modification of Time that allows specification of rules applying to weekdays only, weekends only, or both.
- *Location (Loc)* – allows a user to give access to their location when they are in a pre-specified area. A single rule is defined by a latitude/longitude rectangle.
- *Location and time (Loc/Time)* – allows for rules that are combinations of Loc and Time rules (e.g., "grant access from 3pm-7pm when I am at location A or B.")
- *Location and time with weekends (Loc/Time+)* – modifies Loc/Time to allow rules that apply to weekdays, weekends, or both.

Using the detailed privacy preferences we collected, we calculated the most accurate possible set of rules, or *policy*, for each subject, under each of the different privacy-setting mechanisms. In order to account for differences between subjects, we calculated the policies under varying assumptions about the cost associated with revealing a private location, which we denote as c, and the maximum number of rules users would be willing to specify.

The most accurate policy is the one that reveals as much time possible, while minimizing the amount of time mistakenly revealed. The relative weights assigned to time correctly revealed and mistakenly revealed are determined by c. A small value of c, such as $c = 1$, is equivalent to the assumption that our users would have been equally comfortable with a mistaken reveal as a correct reveal. Larger values of c, such as $c = 100$, are equivalent to the assumption that our users would always err on the safe side, and protect private locations by restricting their policies to never reveal them.

We report the average time shared with advertisers, for each of the privacy-setting mechanisms, under different values of c in Figure 3. Here we see that as the cost of mistakenly revealing a location increases, the policies become more restrictive and the average time shared

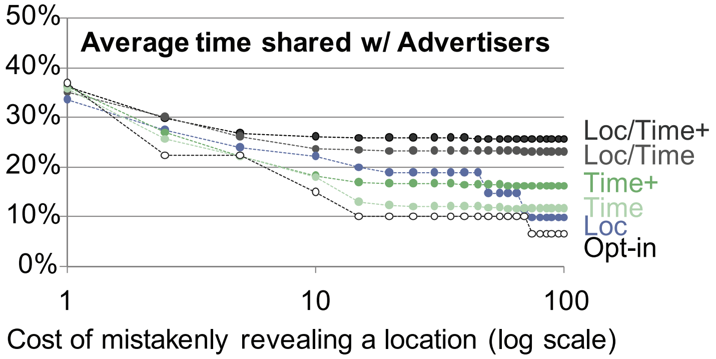

Figure 3: The relationship between percentage of time shared with advertisers under different privacy-setting mechanisms as the cost of mistakenly revealing a location increases.

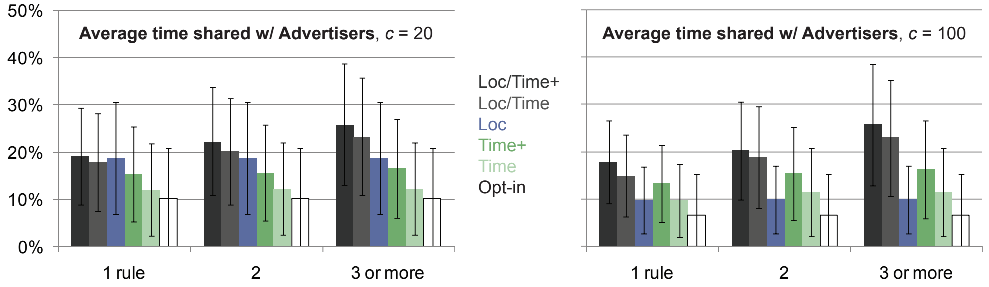

Figure 4: The percentage of time a user would share their location with advertisers under different privacy=setting mechanisms, given a limited number of user-specified rules.

decreases. However, more complex privacy-setting mechanisms, such as Loc/Time+, resist this decrease, and allow policies that maximize the amount of time shared while preventing high-cost mistakes.

For even moderate values of c, such as $c \geq 15$, more complex mechanisms, such as Loc/-Time+ and Loc/Time, result in nearly three times as much sharing as Opt-in, and this difference is statistically significant ($p < 0.05$ for Loc/Time+ and $p < 0.1$ for Loc/Time). This substantial increase in sharing with large values of c is particularly relevant, given that our subjects reported being very concerned about sharing locations marked private with advertisers in our post-study survey.

Additionally, we find that the increases in sharing from more complex setting types can be realized, even if users are only willing to make a small number of rules. As displayed in Figure 4, with $c = 20$, we see a substantial increase in the percentage of time a user would share his or her location with only a single rule. With two rules the differences between the complex mechanisms, Loc/Time+ and Loc/Time, and Opt-in are statistically significant ($p < 0.05$) and marginally significant ($p < 0.1$), respectively. And, as the cost of mistakes increases, the increase in sharing under more complex mechanisms with small numbers of rules is even more dramatic. For example, when $c = 100$ we see an almost three times increase in sharing over Opt-in with a single Loc/Time+ rule, and this increase is statistically significant ($p < 0.05$).

5 Conclusion

Location-based advertising has the potential to fund the growth of future mobile computing systems. Our results suggest that users' strong privacy concerns may hinder this potentially invasive form of advertising, as early efforts reach the market. We also find that advanced privacy settings may help alleviate some of these concerns, making users more comfortable with location-based advertising.

We presented an empirical investigation showing that users' privacy preferences with regards to sharing locations with advertisers are complex. Our survey findings and analysis conducted on over 7,500 hours of location audits, allowed us to characterize these privacy preferences in detail. We find that if users are given only an opt-in/opt-out mechanism, a large percentage will not be able to specify their true privacy preferences, and may simply stop sharing entirely. These findings suggest that designing future systems with more complex privacy settings will benefit all parties, increase sharing, and place more control in the hands of users.

6 Acknowledgments

This work has been supported by a Siebel Scholarship, an NSF IGERT grant, and NSF grants CNS-0627513, CNS-0905562. Additional support has been provided by Nokia, France Telecom, Google, the CMU/Microsoft Center for Computational Thinking, ARO research grant DAAD19-02-1-0389 to Carnegie Mellon University's CyLab, and the CMU/Portugal Information and Communication Technologies Institute. The authors would also like to thank Paul Hankes-Drielsma, Janice Tsai, Tuomas Sandholm, Lucian Cesca, Jialiu Lin, Tony Poor, Eran Toch, and Kami Vaniea for their assistance with our study.

References

[1] L. Barkhuus and A. Dey. Location-based services for mobile telephony: A study of users' privacy concerns. In *INTERACT*, pages 702–712, 2003.

[2] M. Benisch, P.G. Kelley, L.F. Cranor, and N. Sadeh. Capturing location-privacy preferences: Quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing (PUC)*, 2010.

[3] Thorben Burghardt, Erik Buchmann, Jens Müller, and Klemens Böhm. Understanding user preferences and awareness: Privacy mechanisms in location-based services. In *OTM*, pages 304–321, 2009.

[4] Eloise Gratton. M-commerce: The notion of consumer consent in receiving location-based advertising. *Canadian Journal of Law and Technology*, 1(2):59–77, 2002.

[5] Jeffrey Hightower, Anthony LaMarca, and Ian E. Smith. Practical lessons from place lab. *IEEE Pervasive Computing*, 5(3):32–39, 2006.

[6] S. Huang, F. Proulx, and C. Ratti. iFIND: a Peer-to-Peer application for real-time location monitoring on the MIT campus. In *CUPUM 07 - 10th International Conference on Computers in Urban Planning and Urban Management*, July 11-13 2007.

[7] Bernhard Kolmel and Spiros Alexakis. Location based advertising. In *Mobile Business*, 2002.

[8] Clara Mancini, Keerthi Thomas, Yvonne Rogers, Blaine A. Price, Lukazs Jedrzejczyk, Arosha K. Bandara, Adam N. Joinson, and Bashar Nuseibeh. From spaces to places: emerging contexts in mobile privacy. In *Ubicomp*, pages 1–10. ACM, 2009.

[9] Claire Cain Miller and Jenna Wortham. Technology aside, most people still decline to be located. `http://www.nytimes.com/2010/08/30/technology/30location.html?_r=1&scp=1&sq=foursquare&st=cse`.

[10] David H. Nguyen, Alfred Kobsa, and Gillian R. Hayes. An empirical investigation of concerns of everyday tracking and recording technologies. In *UbiComp*, pages 182–191. ACM, 2008.

[11] Bobby L. Rush. Statement by the Honorable Bobby L. Rush, chairman. `http://energycommerce.house.gov/Press_111/20100224/Rush.Statement.2.24.2010.pdf`.

[12] N. Sadeh, J. Hong, L. Cranor, I. Fette, and P. Kelley. Understanding and capturing people's privacy policies in a mobile social networking application. *Journal of Personal and Ubiquitous Computing*, 13(6):1–14, 2009.

[13] Norman Sadeh, Fabien Gandon, and Oh Buyng Kwon. Ambient intelligence: The my-Campus experience. Technical Report CMU-ISRI-05-123, Carnegie Mellon University, July 2005.

[14] J. Tsai, P. Kelley, L. Cranor, and N. Sadeh. Location-sharing technologies: Privacy risks and controls. In *TPRC*, 2009.

[15] Melody M. Tsang, Shu-Chun Ho, and Ting-Peng Liang. Consumer attitudes toward mobile advertising: An empirical study. *International Journal of Electronic Commerce*, 8(3):65–78, 2004.

[16] Ramaprasad Unni and Robert Harmon. Perceived effectiveness of push vs. pull mobile location-based advertising. *Journal of Interactive Advertising*, 7(2), Spring 2007.

[17] Roy Want, Veronica Falcão, and Jonathan Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10:91–102, 1992.