Short Paper

Controlling Location Disclosure by Distinguishing between Public and Private Spaces

Jeremy Wood, PhD, LocationAnonymization.com
Jeremy@LocationAnonymization.com

ABSTRACT
Location sharing apps typically offer users limited options for specifying their complex preferences regarding when and to whom they want their location disclosed. But offering users more extensive options threatens to increase the user burden considerably. We evaluate a method which uses detailed maps to automatically distinguish public and private locations. The results show promise in that the method gives users additional control in disclosing their location data, without significantly increasing the user burden.

Keywords
Location sharing technology, mobile social technology, Privacy

General Terms
Measurement

INTRODUCTION
Typically, an individual's preferences regarding disclosing her location are affected by multiple factors, including where she is, when the information is requested, and the identity of the person making the request [1]. Further complicating matters are the findings that 'people have a hard time articulating effective privacy preferences' [6], and that those preferences are not fixed [7].

A privacy control system that captures users' complex preferences would encourage users to share more. But that goal would be undermined if users were burdened with continually adjusting manifold parameters for multiple contacts [1]. Ideally, we would have a simple system for capturing complex preferences.

This paper addresses that challenge in a manner suggested by [7], who observed that 'users appear more comfortable sharing their presence at locations visited by a large and diverse set of people.' [7] used the crisscrossing of their subjects' GPS tracks to identify such locations (like a university campus), and to distinguish those locations from individuals' homes. [7] referred to these locations as having 'high entropy.'

Distinguishing Public and Private Spaces
We use a different terminology and method. In law and in common sense, there are 'public places' and 'private places,' and people have less 'expectation of privacy' in public [2]. Individuals seem to be particularly concerned to obscure the location of their home; the home is a private place, and people seem especially concerned to keep its specific location private [1].

We use the 'public' and 'private' terminology here. And, we would use this terminology with users because we think that distinguishing between public and private places captures something important about the way that people think about privacy.

With few exceptions, location sharing apps offer users too few options for moderating the granularity of their location disclosure. Typically, apps offer the option to disclose one's precise location, the city in which one is located, or nothing at all. One option which is typically missing and for which there is significant demand is to disclose the 'neighborhood' in which one is located [4].

So, we might suggest giving User X the option: 'Reveal to User Y my precise location if I am in a public place. But if I am in a private place, then reveal my neighborhood/ my current city/nothing.'

METHOD
We tested the method with GPS-derived traces from 21 individuals in the Seattle, WA area. The tracks were centered around King county, and we restricted our study to that county.

In effect, we placed each point on a GIS map obtained from King County, and we determined whether that point was in a public or private location. For purposes of this demonstration, we opted to blur the private locations to show that the method could protect privacy. Of course, it would be possible to blur those points more or to delete them completely, if the user so desired.

Table 1 provides details on how we categorized each point and how we treated coordinates in various locations. We do not blur locations in what we expect to be busy transportation corridors or public destinations (e.g.

shopping malls). And we treat as private every location that is not categorized as public.

Categorized as 'Public'	Categorized as 'Private'
No blurring – GPS coordinates disclosed without alteration • Locations on railroads and all roads more substantial than 'local roads.' • Land use codes that relate to all land uses, except for those that are 'residential.'	Location is blurred to census group block • Everywhere that is not categorized as public

Table 1. Distinguishing Public and Private Locations, and Blurring Locations in Private Spaces.

The GIS programming was fairly straightforward. It was necessary to transform polyline data. And we had to decide how to categorize and treat the various location types. Adjusting these parameters does alter the results. So, there is room for fine-tuning.

RESULTS

Visual inspection of the tracks indicates that the blurring procedure does selectively blur private locations. Specifically, locations in residential areas are blurred to the level of the neighborhood. Images 1-3 show the precise path that individual took to a given residential area, but the individual's locations within that area are automatically blurred.

DISCUSSION

The method does offer a measure of privacy by blurring locations in residential areas. And it would be a simple matter to increase the degree of privacy protection by blurring to a larger area or by dropping completely points in private areas.

The method does have challenges and limitations. Obviously, we would want the method to work reliably, and it can be no more reliable than the map data on which it depends. If there were occasional errors or some such vulnerability, a determined attacker might be able, over time, to collect enough blurred data to pinpoint a location which the user wanted to keep private [3].

Another concern is the fact that an individual's presence in a public area does not assure that the individual is surrounded by a crowd large enough to hide in. So, employed in the manner outlined here, the method may not provide anonymization.

Image 1.

Track 13 -- The individual's locations within this residential neighborhood are blurred to the point in the center. The blurred appearance the pin there indicates that this pin represents many GPS readings.

Image 2.

Track 16 -- The individual's locations within this residential neighborhood are blurred to 4 points. The relative size of the pins indicates that the one in the upper right represents many more GPS readings.

Image 3.

Track 18 -- The individual's locations within a residential neighborhood are blurred to the point represented by the uppermost, right pin. The larger size of that pin indicates that it represents many more GPS readings.

However, the significance of these concerns is limited by the fact that the proposed usage relates to a location sharing application; we can hope that individuals are not choosing to share even blurred location tracking information with determined attackers. But this danger reminds us that location sharing involves some inherent risks, regardless of the software.

CONCLUSION

The method demonstrated here distinguishes public and private places automatically. Thus, users would be spared the burden of geofencing their homes and other sensitive areas, and of updating their preferences each time they change location. And users could be offered the option 'Reveal to User Y my precise location if I am in a public place. But if I am in a private place, then reveal nothing/my neighborhood/ my current city.' In this manner, the method could help users articulate, communicate, and effectuate their location privacy preferences.

REFERENCES

1. Benisch, M., Kelley, P. G., Sadeh, N., and Cranor, L. F. Capturing Location-Privacy Preferences: Quantifying Accuracy and User-Burden Tradeoffs. (2010)
2. Katz v. United States 389 U.S. 347, (1967)
3. Krumm, J. Inference attacks on location tracks. In PERVASIVE'07: Proceedings of the 5th international conference on Pervasive computing. Springer-Verlag, (2007), 127-143.
4. Lin, J., Benisch, M., Sadeh, N., Niu, J., Hong, J. Lu, B., and Guo, S. A Comparative Study of Location-sharing Privacy Preferences in the U.S. and China (2012)
5. Palen, L., and Dourish, P. Unpacking "Privacy" for a Networked World. *ACM. (*2003)
6. Sadeh, N., Hong, J., Cranor, L., Fette, I. Prabaker, M. and Rao, J. Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application. Human-Computer Interaction Institute. (2007)
7. Toch, E., Cranshaw, J., Drielsma, P. H., Tsai, J. Y., Kelley, P. G., Springfield, J., Cranor, L., Hong, J., and Sadeh, N. Empirical Models of Privacy in Location Sharing. UbiComp, ACM. (2010)
8. J. Wood. Method of Providing Location-Based Information from Portable Devices. United States Patent 8,185,131. (2012)

ACKNOWLEDGEMENTS

Special thanks to Christian Frugard, John Krumm, and David Lazer.