Location Disclosure

Alex Endert
aendert@cs.vt.edu

Location Disclosure Overview

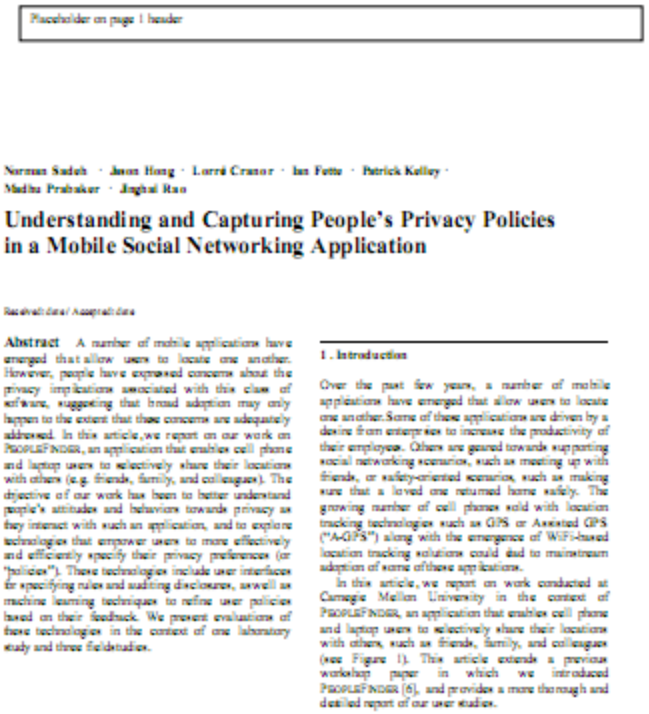

Placeholder on page 1 header

Norman Sadeh · Jason Hong · Lorrie Cranor · Ian Fette · Patrick Kelley · Madhu Prabaker · Jinghai Rao

Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application

Received date / Accepted date

Abstract A number of mobile applications have emerged that allow users to locate one another. However, people have expressed concerns about the privacy implications associated with this class of software, suggesting that broad adoption may only happen to the extent that these concerns are adequately addressed. In this article, we report on our work on PEOPLEFINDER, an application that enables cell phone and laptop users to selectively share their locations with others (e.g. friends, family, and colleagues). The objective of our work has been to better understand people's attitudes and behaviors towards privacy as they interact with such an application, and to explore technologies that empower users to more effectively and efficiently specify their privacy preferences (or "policies"). These technologies include user interfaces for specifying rules and auditing disclosures, as well as machine learning techniques to refine user policies based on their feedback. We present evaluations of these technologies in the context of one laboratory study and three field studies.

1. Introduction

Over the past few years, a number of mobile applications have emerged that allow users to locate one another. Some of these applications are driven by a desire from enterprises to increase the productivity of their employees. Others are geared towards supporting social networking scenarios, such as meeting up with friends, or safety-oriented scenarios, such as making sure that a loved one returned home safely. The growing number of cell phones sold with location tracking technologies such as GPS or Assisted GPS ("A-GPS") along with the emergence of WiFi-based location tracking solutions could aid to mainstream adoption of some of these applications.

In this article, we report on work conducted at Carnegie Mellon University in the context of PEOPLEFINDER, an application that enables cell phone and laptop users to selectively share their locations with others, such as friends, family, and colleagues (see Figure 1). This article extends a previous workshop paper in which we introduced PEOPLEFINDER [6], and provides a more thorough and detailed report of our user studies.

Norman Sadeh
ISR - School of Computer Science - Carnegie Mellon University
5000 Forbes Avenue - Pittsburgh PA 15213-3891
sadeh@cs.cmu.edu

Developing Privacy Guidelines for Social Location Disclosure Applications and Services

Giovanni Iachello
College of Computing and GVU Center
Georgia Institute of Technology, USA
giac@cc.gatech.edu

Ian Smith, Sunny Consolvo, Mike Chen
Intel Research
Seattle, WA, USA
{ian.e.smith, sunny.consolvo, mike.y.chen}@intel.com

Gregory D. Abowd
College of Computing and GVU Center
Georgia Institute of Technology, USA
abowd@cc.gatech.edu

ABSTRACT

In this article, we describe the design process of Reno, a location-enhanced, mobile coordination tool and person finder. The design process included three field experiments: a formative Experience Sampling Method (ESM) study, a pilot deployment and an extended user study. These studies were targeted at the significant personal security, privacy and data protection concerns caused by this application. We distil this experience into a small set of guidelines for designers of social mobile applications and show how these guidelines can be applied to a different application, called Relo. These guidelines cover issues pertaining to personal boundary definition, control, deception and denial, and group vs. individual communication. We also report on lessons learned from our evaluation experience, which might help practitioners in designing novel mobile applications, including the choice and characterization of users for testing security and privacy features of designs, the length of learning curves and their effect on evaluation and the impact of peculiar deployment circumstances on the results of these finely tuned user studies.

Categories and Subject Descriptors

D2.2 [Design Tools and Techniques]: Evolutionary Prototyping; D.2.1 [Software Engineering]: Requirements/Specifications—elicitation methods; K.4.2 [Computers and Society]: Social issues; K.8.m [Personal Computing]: Miscellaneous

General Terms

Design, Human Factors, Security

Keywords

Social Mobile Applications, People Finder, Iterative Design, Prototyping, Privacy

1. INTRODUCTION

We are interested in designing social mobile applications, i.e., mobile Information Technology (IT) applications which facilitate everyday social interactions. Social mobile applications include text messaging (texting) services, person finders, and availability managers. During the last few years, a slew of specialized applications has emerged, thanks to the availability of powerful computing platforms (the most common being smart phones and networked PDAs), infrastructure software and novel context sensing techniques. In the present article, we will concern ourselves specifically with social location disclosure applications, that is, applications that enable the communication of location among individuals within their social networks.

These applications are widely considered to have a strong commercial potential, especially with consumers in younger age groups. For example, a market survey of US cell phone users conducted in 2004 showed that person finder applications are the second most popular choice among data-intensive applications people would use on their cell phones if they were to spend an additional USD 5–10 on their monthly bill [21].

Despite their promising commercial outlook, the applications that have been launched to date have not performed well in the marketplace. The most widely deployed person finder in the US mobile phone market, AT&T Find People Nearby, has failed to become a large market success, and AT&T Wireless's new parent company (the operator Cingular) may discontinue the application as part of the merger [2]. Other person finders, such as Dodgeball [6], which do not rely on operator support, have a fringe following of dedicated users, but are far from being widespread. Child tracking applications, corporate employee management and similar specialized services, deployed in several countries (e.g., UK, Japan) in collaboration with operators have experienced somewhat better success.

The reasons for the lukewarm acceptance of social location disclosure applications may lay in several interrelated fac-

PeopleFinder Paper, Meet the Authors

Norman Sadeh

Professor, CMU

Jason Hong

Assistant Prof., CMU

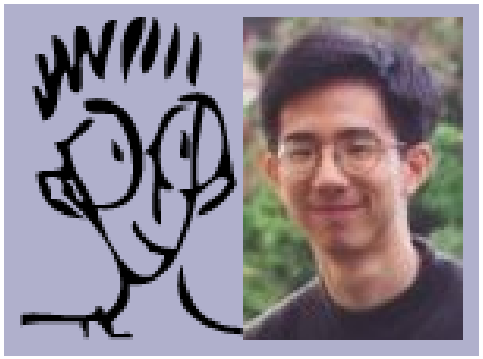

Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, Jinghai Rao

What is this paper about?

- Using *PeopleFinder* "to better understand people's attitudes and behaviors towards privacy"

- Presentation of evaluations of technology
  - 1 lab study
  - 3 field studies

PeopleFinder

- Invited users can see your location
  - Based on user defined policies
- Location
  - GPS
  - GSM Triangulation
  - WiFi Location

Figure 1. PEOPLEFINDER is an application that lets users share their locations with others subject to privacy policies they can refine over time.

PeopleFinder

- PEA = Policy Enforcing Agent
- Location updated regularly, uses “last seen”

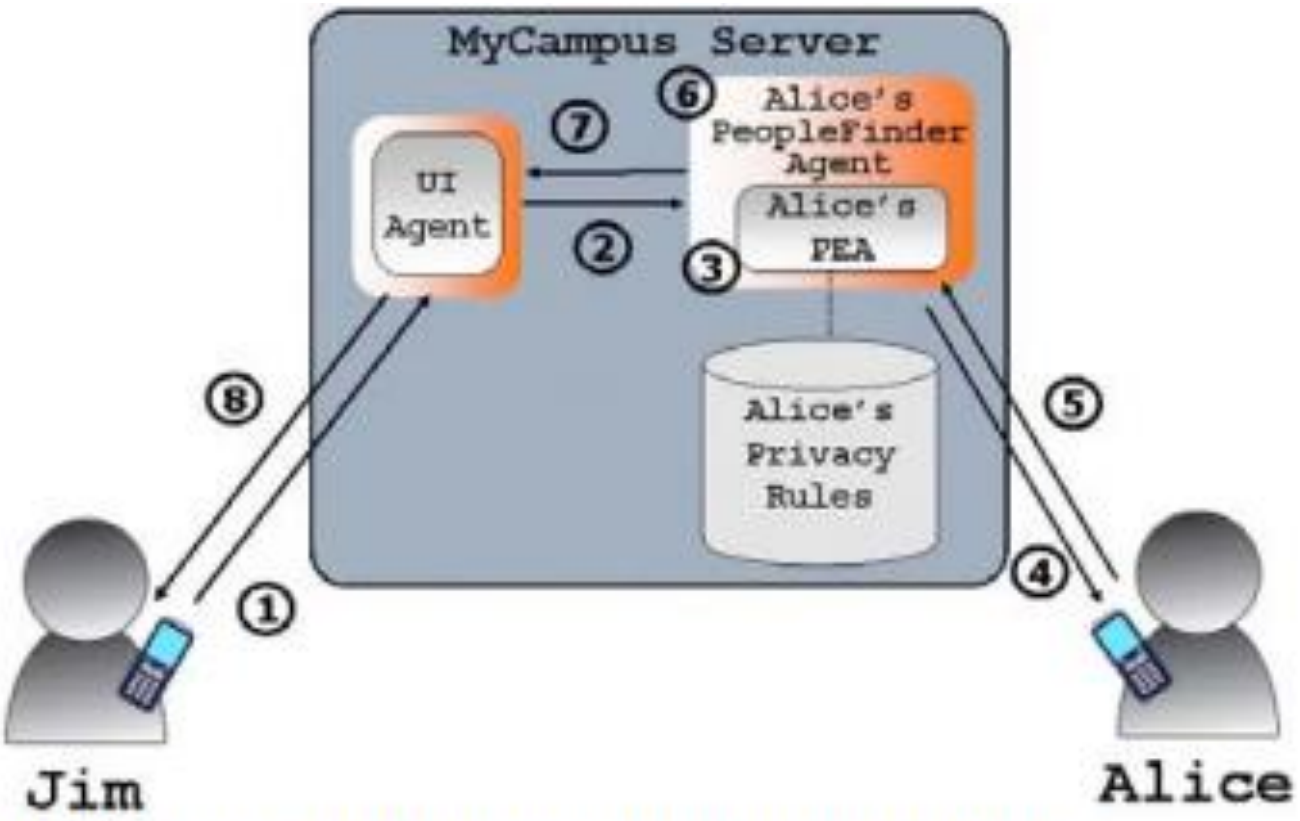

Figure 3. Processing Jim's request for Alice's location.

“Requesting User” “Target User”

PeopleFinder

- Interface to find friends

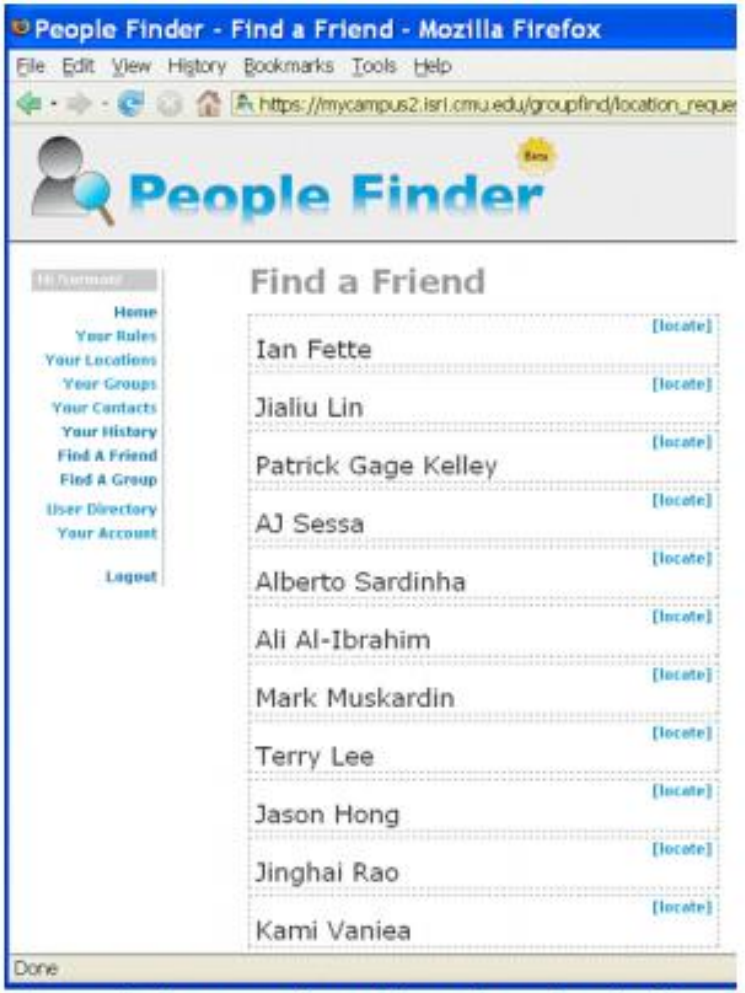

Figure 2. Browser-based interface for finding the location of a person. Equivalent Java and C# applications are also available for laptops and several cell phones.

PeopleFinder

- Interface to create rules
  - Denying a user sends “ambiguous” return message

Figure 4. User interface for defining simple privacy rules.

PeopleFinder

- Define blocked areas
  - Falls back to “last seen” location

Figure 5. Users can also define locations as combinations of rectangular areas for use in location-sensitive privacy rules.

PeopleFinder

- Runs on Windows Mobile, Windows and Mac laptops

- Found notifications “make users feel more comfortable with application”

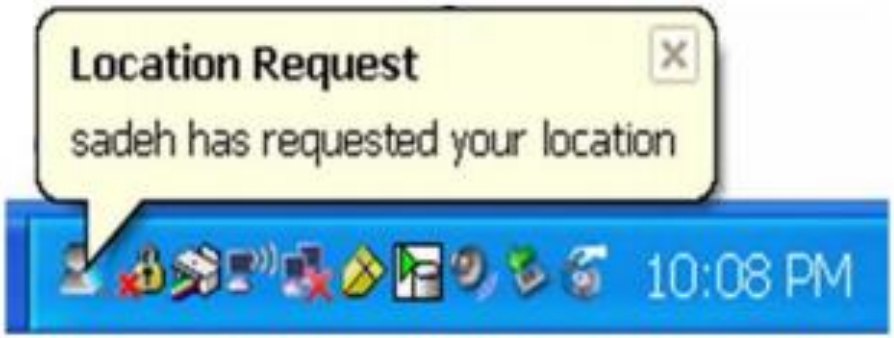

Figure 6. Bubbles notifying users of incoming queries help maintain awareness while being minimally disruptive.

User Feedback and Auditing

Figure 7. Auditing functionality helps users understand how their policies work and enables them to more effectively refine their policies.

Figure 8. Explanation can help users better understand their policies. User feedback can also be used to make suggestions or learn the user's preferences.

Lab Study

- 19 participants
  - Later study with 60 participants
- Asked to disclose information such as:
  - “My colleagues can only see my location on weekdays and only between 8am and 6pm”
- Created 30 individual scenarios

Lab Study Findings

- Specifying initial rules ~5 minutes
  - ~8 minutes if the user modified rules on the fly during study
- Initial rules correctness

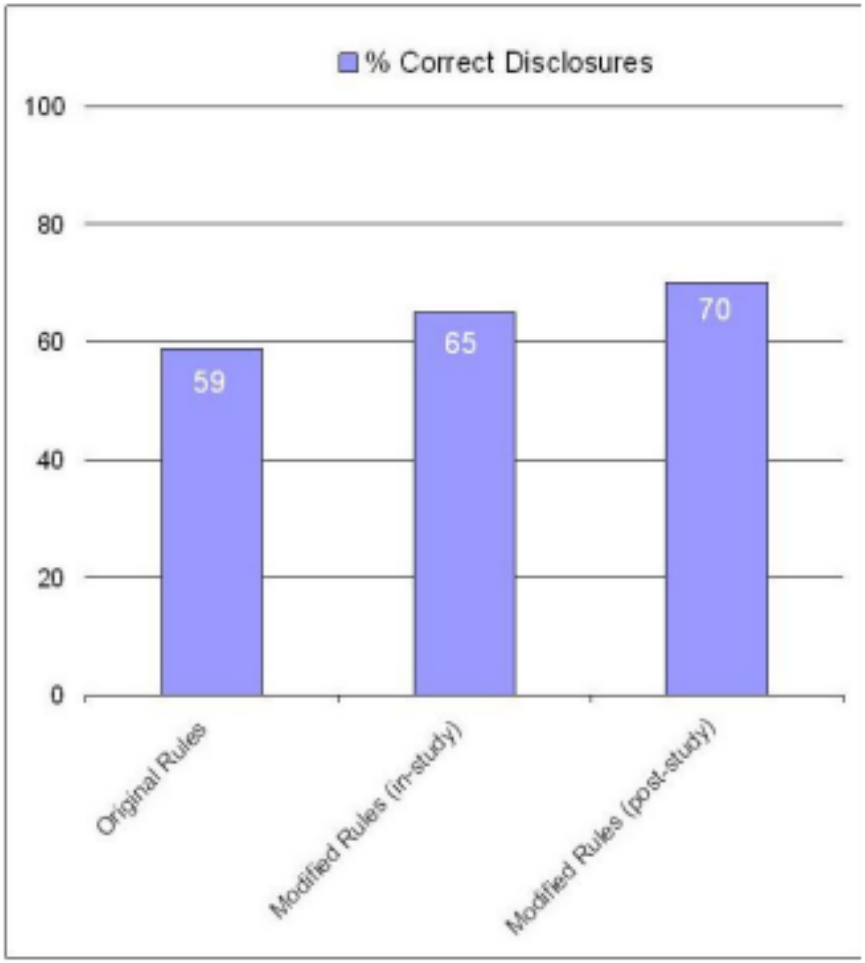

Figure 9. Controlled lab experiments: Users are not very good at articulating their privacy policies – accuracy of initial rules versus rules modified after being presented with 30 customized usage scenarios.

Lab Study Findings: How Many Rules?

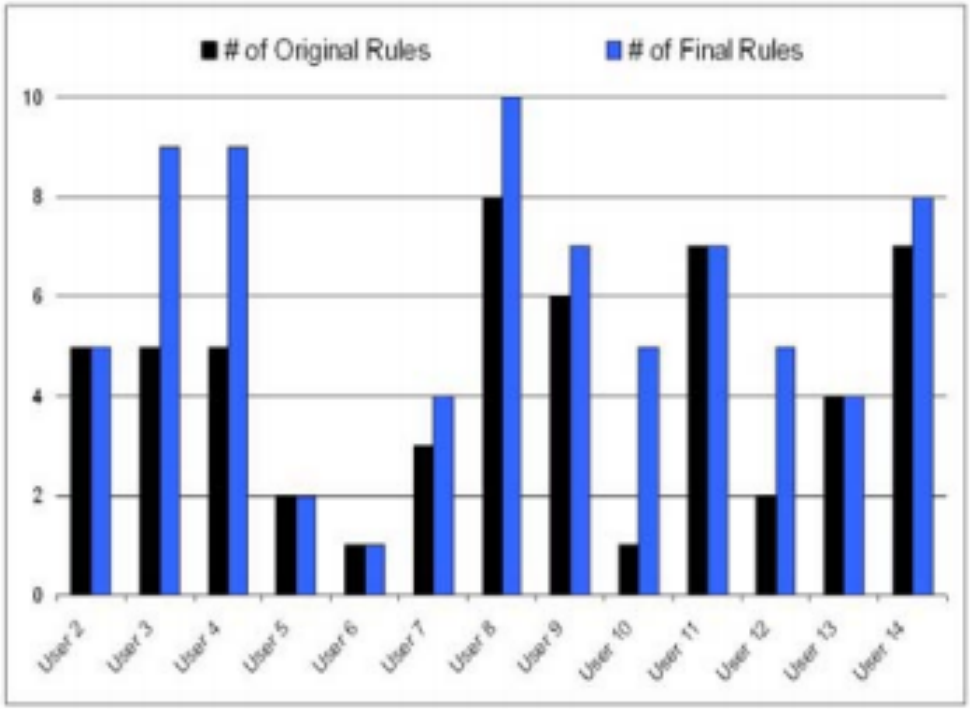

Figure 10a. Controlled lab experiments: initial number of rules versus final number of rules. User 1 was used for a pilot study and thus is not included in these results.

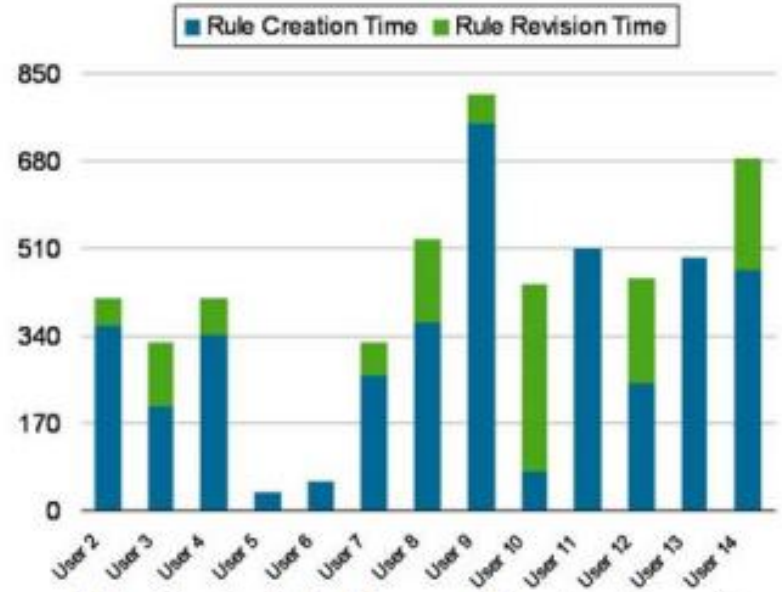

Figure 10b. Controlled lab experiments: time (in seconds) spent creating and modifying rules – the latter includes both changes to initial rules and addition of new rules.

Lab Study Findings

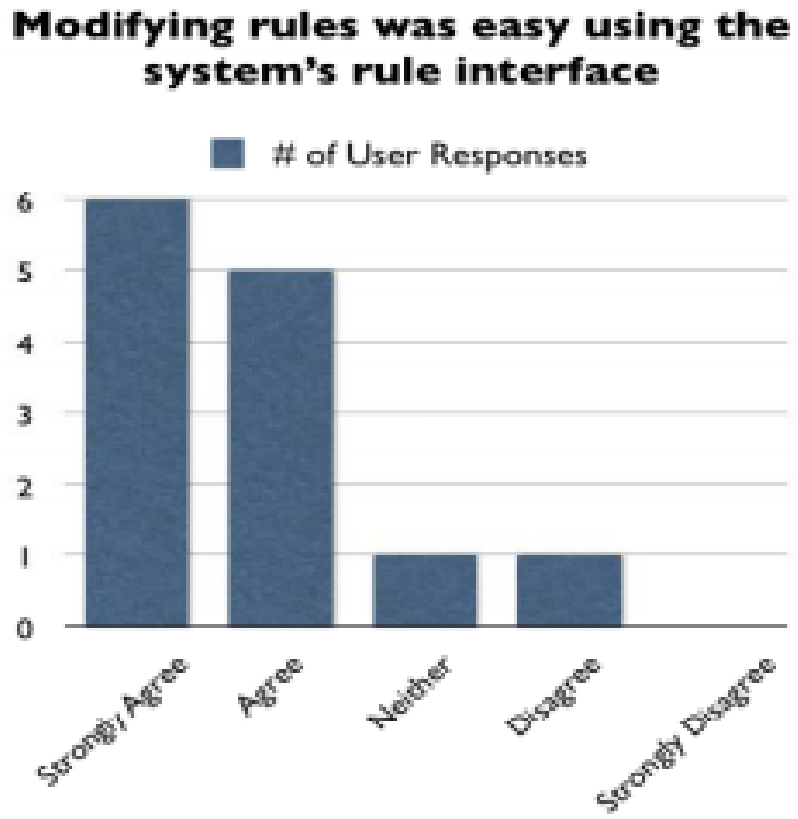

Figure 11. Controlled lab experiments: user feedback suggests that difficulties in articulating policies are not due to a poorly designed rule interface.

Weird way to phrase a question?

Lab Study Findings

- Little correlation between:
  - Policy accuracy and number of specified rules
  - Policy accuracy and time spent defining/refining rules
- Users “reach a plateau and are often unable to articulate highly accurate policies”

PeopleFinder Field Study

- 3 “in the wild” groups:
  1. Their Research Group (15 users)
  2. MBA Students (7 users)
  3. Buggy Race Organizers (6 users)

Field Study Findings

- Machine Learning algorithms show promise
  - Based on user feedback
- Rules show 79% accuracy

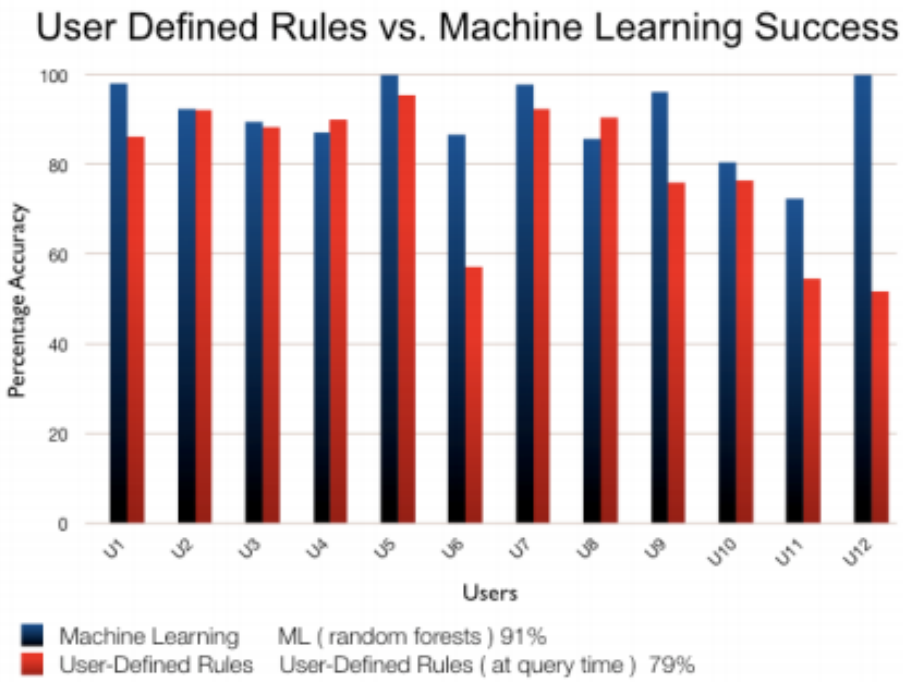

Figure 15. Field studies: accuracy for 12 most active target-users from 3 field pilots involving over 60 users. A random forest classifier shows promise in helping improve the accuracy of user-defined policies.

Overall, …

- Short initial setup time gives 65-79% accuracy, with some rules developing over time
  - Allow user to pick pre-defined patterns?
- Can machine learning for rules help? They think so…
  - Users are not very effective specifying highly accurate rules
- Blacklist (information is disclosed unless specified) vs. whitelist (only disclosed if specified).
  - Manageability vs. privacy
- Users “relax” with the release of location with time

Reno and Boise Paper

- Ga. Tech
  - Giovanni Iachello
  - Gregory Abowd
- Intel Research, Seattle
  - Ian Smith
  - Sunny Consolvo
  - Mike Chen

What is this paper about?

- Workshop outcome, agenda?
- “developing privacy-observant application that allows people to communicate their location”
- Three studies:
  - Experience Sampling Method Study
  - Pilot Deployment
  - Extended User Study
- Discovered guidelines for social mobile developers

First, understand the user!

- 2 week study, 16 adults
  - *What* are people willing to disclose about location
  - Diary study
  - Users interrupted randomly throughout day with hypothetical location request.
- Results:
  - Either disclose only useful location info OR deny
  - No *blurring* discovered
    - Intentionally vague
  - *Who* is requesting, *why* do they need to know, *what* would be the least amount useful, *am I willing to share that?*
    - Want to *stretch the truth!*
  - *"Okayness checking"* – did you make it home ok?

Reno

- Mobile app to disclose location
  - Ability to “learn names” of locations
    - E.g. “Home”, not address
- Ignoring = deny
- Nearby locations to choose from
- Pilot Deployment (8 users)

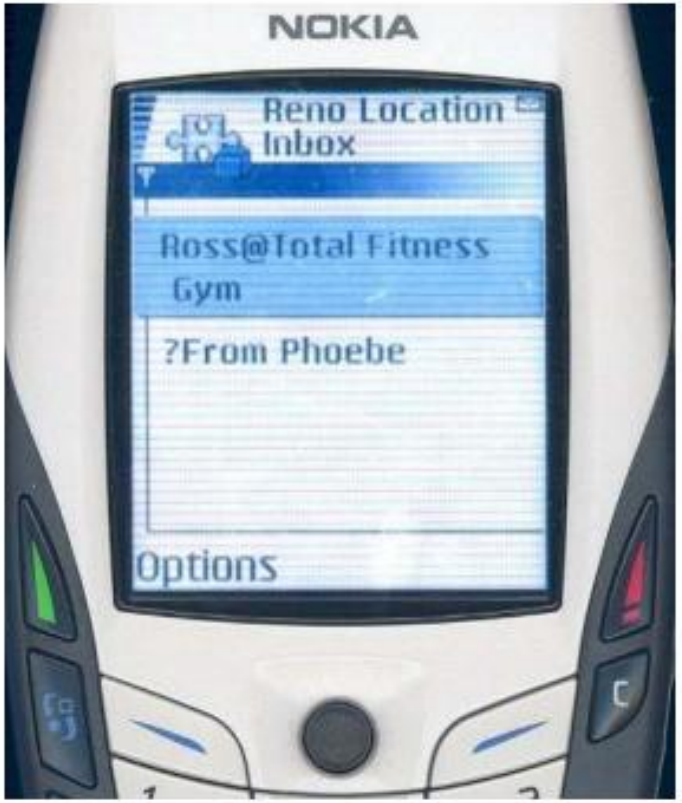

Figure 1: Screenshot of The Reno Application. The location “inbox” shown here contains two messages, a disclosure from Ross and a query for the user’s location from Phoebe.

Pilot User Feedback

- Recipient of location used knowledge to further investigate issue:
  - “I’m at bus stop”
    - Using time of day, day of week, usual schedule, etc.
    - Means you will be home in 15 minutes

Study #2

- 2 week study
  - “modified” version of Reno
  - “waypoints” for *instant reply list*
- Results:
  - “location” may not be what they want to disclose.
    - E.g. “on the way home”, not location
  - Automatic response was not liked

Design Guidelines

- Automatic Reply should be feature, not default
- Users choose reply to location request
- Support ignoring requests
- Deceiving replies support
- Ability to signal “busy” / away messages
- Person-to-person communication before group
- Do not use centralized services

Outcome: Boise

- Features 3 modes:
- Normal
  - Query other users on location
- Tracking
  - Allow select user to “track” you.
- Away
  - Provides automatic away message when location requested.

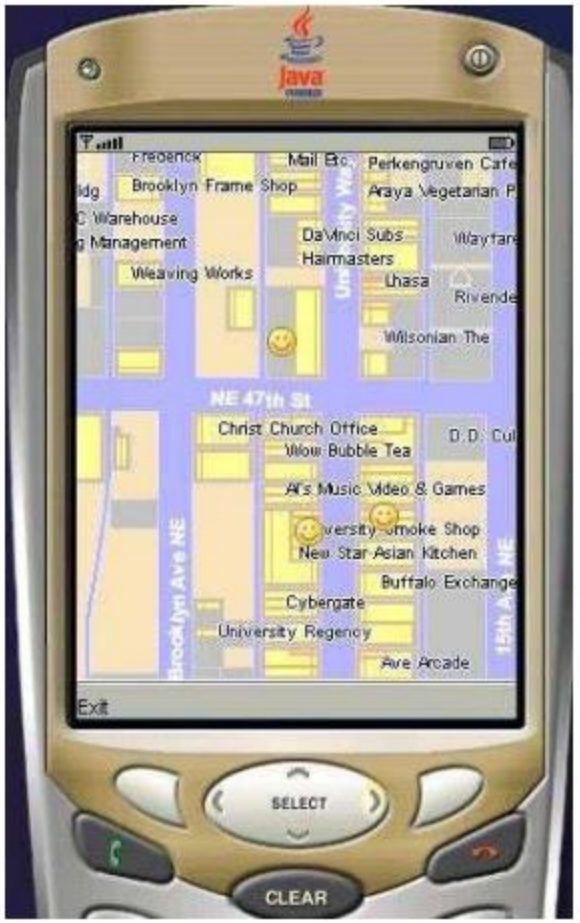

Figure 2: Prototype implementation of Boise. The owner of the device is located at the center of the image (at the *Wow Bubble Tea Shop* in Seattle's university district). Three other people (the smiling faces) have disclosed their location to the user. The background map was taken from Lost In Seattle. (www.lostinseattle.com)

Discussion

Discussion

- User Location Deception?
- What location information response?
  - Location? Where I'm going? How long until I get there?
- "Blocked Areas", do they work?
- Do users know what rules/policies they want?
- Focusing on only most active users, good design?