



Learning from Context: A Field Study of Privacy Awareness System for Mobile Devices

L Jedrzejczyk and C Mancini and D Corapi and B A Price and A K
Bandara and B Nuseibeh

09 June, 2011

Department of Computing
Faculty of Mathematics, Computing and Technology
The Open University

Walton Hall, Milton Keynes, MK7 6AA
United Kingdom

<http://computing.open.ac.uk>

Learning from Context: A Field Study of Privacy Awareness System for Mobile Devices

L. Jedrzejczyk*, C. Mancini*, D. Corapi[§], B. A. Price*, A. K. Bandara*, B. Nuseibeh*[†]

*Department of Computing, The Open University, Milton Keynes, UK

[§] Imperial College London, UK

[†]Lero, University of Limerick, Republic of Ireland

{L.Jedrzejczyk; C.Mancini; B.A.Price; A.K.Bandara; B.Nuseibeh}@open.ac.uk, d.corapi@imperial.ac.uk

ABSTRACT

In this paper we investigate the effectiveness of context-awareness and machine learning in ensuring social acceptance of real-time feedback in a social location tracking system. Real-time feedback is a novel privacy feature supporting bi-directional function of privacy management. Its main function is to deliver feedback to the user by using an appropriate form of notification every time a user's location has been checked. We evaluated our technology in the context of Buddy Tracker, our context-aware, location-sharing application for mobile devices. We report on our experience from the development process and also discuss findings of our field study with 15 participants. The findings show that context-awareness and machine learning can minimize the intrusiveness of real-time feedback, therefore making this important function socially more acceptable thus allowing users to benefit from the increased level of awareness that real-time feedback affords. We conclude with recommendations on how a better understanding of the user and application-specific context can improve the user experience and social acceptance of the system.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: Evaluation/methodology, Graphical user interfaces (GUI), User-centered design. H.3.4 [Systems and Software]: Current awareness systems, user profiles and alert services.

General Terms

Design, Experimentation, Human Factors, Theory.

Keywords

Feedback, mobile computing, location based services, privacy management, social translucence, context-awareness.

1. INTRODUCTION

According to Altman [1], privacy in the physical environment can be regarded as an ongoing process of regulating boundaries. At the heart of Altman's theory is an environment that provides tools and mechanisms for regulating privacy [2]. The environment can be regarded as physical structures (walls, position of physical items) [3], [4], systems [4], meaningful places, social actions, or events [5] that determine our behavior [1]. A key property of Altman's privacy regulation theory is *bi-directionality*, whereby privacy regulation is a social process involving input from others (i.e. noise, previous experience) and output to others (e.g., communication).

Longitudinal accumulation of experience with an environment builds social awareness, i.e. a shared knowledge that helps us structure our interactions with one another. According to Dourish

and Bellotti *awareness is an understanding of the activities of others, which provides a context for your own activity* [6]. We argue that awareness is an important element of the privacy management process, because it conditions our social interaction [7], affects our privacy decisions [8] and impacts our comfort in sharing information [9]. In this respect, the task of a privacy-aware system designer is to create environments (system architectures) that can incorporate artifacts from the physical world into digital systems and influence user behavior in a way that supports continual privacy management [10].

Continuous boundary regulation in both online and offline activities requires people's ability to share their information with others but to also sense input from others (e.g., when others make comments about them, respond to them during a conversation or view who looked their location information). In the physical world, when someone walks in the street they can be seen by others but can also see others around them. However, apart from academic applications [11], location-sharing applications make it easy to share but not easy to sense [12], in that the user can be seen without their awareness. This contradicts Altman's bi-directional property of privacy, which says that privacy regulation requires both output (sharing) and input (sensing, feedback). But while in the real world it is easy to absorb and use information from the environment, accounting for Altman's bi-directionality of privacy into digital systems poses a technological challenge. Our main objective in designing a real-time feedback feature for location-sharing applications was to create a technology that could support bi-direction and thus (similarly to [13]) help users understand each other's actions with respect to their privacy and social relationships.

We borrowed from Erickson and Kellogg's concept of social translucence [7] in supporting awareness through shared understanding that enforces accountability by making things visible to one another. We did so by integrating a real-time feedback feature in our custom-built location-sharing application: every time a user checks another user's location, the system logs the request and automatically sends a notification to the data owner. This means that the data owner is informed of every check made on their location.

Based on our previous work on such feedback mechanisms [8], we found that real-time feedback technology has the potential to improve users' awareness and protect privacy. However, we also found that this technology raises a number of social and practical issues, and identified several conditions that need to be met if such technology is to become socially accepted. We found that the acceptability of any feedback notification form depends on the particular context of the receiver at any given time: e.g., an

audible message could be appropriate while driving a car but not during a business meeting. Notification types that are inappropriate for a particular context (for example, because of their intrusiveness or due to poor timing) could result in the system being rejected by users altogether. Therefore it is important that the system be capable of sensing user’s context and selecting an appropriate feedback representation for any given situation.

In this paper we share our experience of building and evaluating context-aware real-time feedback in *Buddy Tracker*, our custom-built mobile location-sharing application. Our system is capable of sensing the environment thereby supporting the continual and selective disclosure of personal information by providing real-time feedback as a way of informing users about how their location information is used. We show that the combination of context-aware real-time feedback coupled with machine learning techniques can reduce the intrusiveness of technologies affording social translucence. Our findings also suggest that appropriateness of notifications contributes towards greater trust in the system, greater level of comfort and enhanced user experience.

The next section discusses previous work in the field of HCI, relevant to the concept of feedback. This is followed by an overview of the Buddy Tracker system, which we developed for the purpose of the study, and a discussion of the system’s technical details, including its architecture, context-awareness mechanism and learning engine. We then present the findings of our study, including a survey and a field trial of context-aware real-time feedback use. In the final section we summarize our findings, drawing our conclusions.

2. RELATED WORK

Bellotti and Selen [14], described feedback as “*informing people when and what information about them is being captured and to whom the information is being made available*”. HCI researchers generally consider this as important privacy feature for context-aware applications, which has been discussed extensively in prior work [8], [9], [14-16]. Feedback has been recognized as an important factor allaying users’ privacy concerns and increasing comfort of sharing location [9]. Research has also shown how feedback information can affect users’ behavior: for example, in previous work [8] we demonstrated how feedback was used to enforce social norms by making people more accountable for their actions while using location-sharing application.

Beyond examining the social consequences of using feedback, a number of researchers have proposed design solutions for representing feedback in digital systems. One of the first attempts to design unobtrusive real-time notifications for media spaces suggested that a LED light was a sufficient and non-intrusive way of informing people about data collection [14]. Sellen et al. proposed a novel design for a situated device called *The Whereabouts Clock*, which informed group members about the current location and activities of others [17]. Access notifications proposed by Hong provided both feedback and control, which allowed users to react to location requests while interacting with the system [18]. Feedback design also considers information to be presented to users in the form of disclosure logs: for example, Lederer [19], Tsai [9] and Sadeh [20] designed interfaces for presenting a history of location requests for large screen devices. Additionally, our previous work [18] and that of Raento and Oulasvirta [8], [15] discuss the concept of historical feedback for mobile applications.

A number of design considerations for real-time feedback were proposed by Hong [18], Hsieh [16] and Sadeh [20] for large screen devices. However, we are not aware of any work exploring feedback design in the context of mobile devices. These are clearly very different from traditional, large screen computers due to the diversity of their context of use and their ubiquitous nature. Further, the work described so far has strongly focused on visual feedback, which, as our findings show, is not always appropriate in a mobile context. Hence our interest is in designing a real-time feedback system that will support awareness, an important element of the bi-directionality of privacy, while also ensuring the effectiveness but also unobtrusiveness of notifications.

Nguyen and Mynatt highlighted the danger, in invisible (ubiquitous) computing, of “*interfaces that do not give people the needed tools of awareness and control to understand and shape the behavior of the system*” [21]. In this respect, our goal is to design *invisible* notification systems that will be well integrated into the life style, practices and values of each individual user. By *invisible* we mean, the state of technological unobtrusiveness, in which people can absorb feedback information without cognitive effort. In our design considerations we explore several sensory dimensions for representing feedback to extend the traditional way of communication between the mobile device and the user. We also use the notion of context-awareness to enrich the user experience and provide support for *contextual-adaptation* [22].

3. BUDDY TRACKER

Buddy Tracker is a mobile phone application that allows users to share their location amongst a group of people. The application comprises two main components, a client application developed in Java for Android devices and a server, which has been implemented as a set of PHP classes working on top of the MySQL database. HTTP is used for communication between the client and the server. Encryption mechanisms are used to prevent eavesdropping of the sensitive contextual information transmitted from the client device. The architecture of the Buddy Tracker system is illustrated in Figure 1.

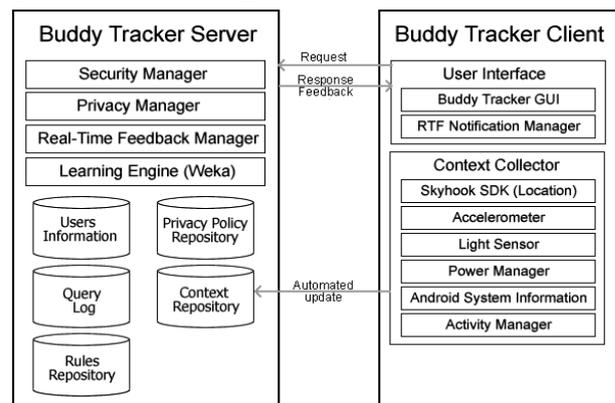


Figure 1. The Buddy Tracker architecture

3.1 Buddy Tracker Server

The server implements four modules (Security Manager, Privacy Manager, Real-Time Feedback Manager and the Learning Engine), and uses five main data repositories (User Information, Privacy Policy Repository, Query Log, Context Repository and

Rules Repository). The User Information repository contains information about users, such as their name, login and password. The Context Repository stores the users' location information and other contextual data (described in section 3.4). Users' privacy preferences are stored in a Privacy Policy Repository and the Query Log contains information about location requests. This repository is used by Buddy Tracker's aggregated feedback module, enabling users to view who accessed their location in the past. The remaining element, the Rules Repository, contains information about users' preferences for real-time feedback.

The functionality of Buddy Tracker's server modules can be illustrated using an example scenario in which a user looks up the location of another. The first module that takes part in that request is the Security Manager, which is responsible for each user's authentication. After a successful check of a user's details in the Users Information repository, the location query is forwarded to the Privacy Policy Repository which analyzes the data owner's privacy policy. The system sends a response to the user based on the requester's details and the data owner's privacy policy. Information about the location query (data requester, data owner, location, granularity level of disclosed location) is then forwarded to the Real-Time Feedback Manager. The Real-Time Feedback Manager first checks the data owner's preferences for real-time feedback (stored in Rules Repository) and then sends the feedback notification based on that information. Secondly, the Real-Time Feedback Manager saves the location request information in a Query Log for future reference.

3.2 Buddy Tracker Client

The Buddy Tracker client application is implemented as an Android application. The client application consists of two elements: the user interface, which allows users to control the disclosure of their location, change privacy preferences and check their buddies; and the background service, which automatically updates the user's current context and checks for new location lookups.

3.2.1 Functionality

The primary interface of Buddy Tracker shows five tabs:

1. Home – main screen, allowing the user to control the frequency of location updates and switch on/off the background service.
2. Buddies - shows the list of all friends with a link to the map at the bottom of the list, which shows all the user's friends on a single map. Clicking on a friend's name opens their profile, with more detailed information about current location as a text description with a link. In this section the user can also adjust their privacy preferences (i.e. hide/blur their location for chosen person, or hide their location for a specified amount of time).
3. Settings – allows the user to adjust their quiet hours (time, when the system should not deliver notifications, e.g. the user can tell the system that they do not wish to receive any notifications between 10PM and 7AM) and enable preferred representations of the real-time feedback.
4. Feedback - aggregated feedback mechanism, which allows users to see who has viewed their profile, location or who accessed their location history. Users could also access aggregated feedback information directly from the buddy profile.
5. Forms – displays a list of recent ESM questionnaires awaiting user's feedback. ESM questionnaires are described in detail in section 4.2.2.

3.2.2 Sensory Dimension of Real-Time Feedback

The sensory dimension relates to the feedback representation, describing how information will be communicated to users. We have identified three subgroups of the sensory dimension:

- **Auditory feedback** - describes any audio interaction between the system and the user, which has been recognized as an intuitive and unobtrusive medium for communication. It can be as simple as a distinct musical tone playing when the event occurs or it can incorporate fully descriptive natural language feedback.
- **Visual feedback** - relates to any visual element or feature on a mobile device that supports interaction including GUI elements used in ad-hoc communication. It can be used to represent the current state of the system, and also to display aggregated information based on historical data, i.e. icons, warnings, dialog boxes (Figure 2a), privacy critics [23], disclosure logs [19], map visualizations (Figure 2c) or status bar notifications (Figure 2b). Visual feedback can be also represented via hardware features, which relate to any visual feature of mobile device design that can be managed programmatically and used for communication (e.g. the LED light in HTC Desire Android phone). Selected interfaces for visual feedback are presented in Figure 2.
- **Tactile feedback** - describes the vibro-tactile interaction between the system and the user such as the phone vibrating when an event occurs.

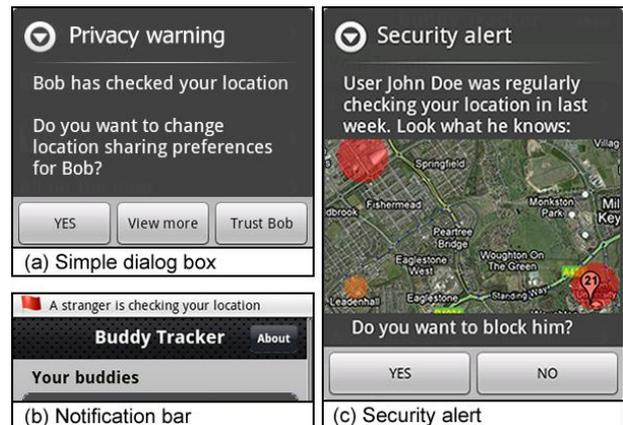


Figure 2. Selected visual representations of real-time feedback interfaces.

3.3 Context in Buddy Tracker

Results from our previous studies [8] on the efficacy of real-time feedback suggest that the most important criterion for the social acceptance of real-time feedback is the ability to convey meaningful information in the most appropriate way for the given situation. It requires the mobile application's ability to sense the environment (*contextual sensing* [22]) and adapt the user interface to the given situation (*contextual adaptation* [22]).

Table 1. Types, representations and descriptions of contextual information collected by Context-Manager component.

Type	Rep.	Description
Identity	User information	Information about the user, age, gender
Time	Timestamp	Date and time of update
Location	coordinates	Collected through GPS, WiFi or nearest cell-id
Location	Full Address	Collected through Google Geo Service API
Location	Category	Category of location, i.e. work, gym, shopping centre, library. We used categories from [24]
Location	Personal tag	User defined description of the location, i.e. home, John's
Screen state	0,1	Says if the screen was on/off
Current task	Package name	Currently used application, i.e. com.android.camera suggests user was using a camera
Phone mode	Normal, In call or Ringtone	Describe the current state of the phone
Ringer mode	Silent, Vibrate or Normal	Current ringer setting
Battery level	0-100	Battery level shown as percentage
Battery charging	0,1	Tells if battery was plugged to the charger
Light level	Number	Current light level in lux
Light description	i.e. dark, bright inside	Current light level expressed in natural language
Phone in use	0,1	Says if the user was using the phone
Movement	x,y,z	Reading from the accelerometer
Phone position	Number	Current position of the phone, calculated based on the current accelerometer reading.
Screen visibility	Text	Textual representation describing probability of the screen's visibility: unknown, visible, maybe_visible or invisible
Screen brightness	0-255	Numeric representation of the screen brightness
Company	0,1	Tells if user is likely with others. Collected from users' calendar, i.e. meeting entry suggests not alone
Company relationship	Text	Label representing relation between the user and the company, we used categories from [24]
Current activity	Text	Real world activity, collected from the calendar or inferred through the current location

In order to use the user's context effectively, we started our work on context-aware real-time feedback with exploratory studies

aimed at understanding what 'mobile context' means for our application, and how we can utilize that. We conducted interviews and organized a focus group with users of mobile devices to understand how contextual factors have an impact on users' preferences for the real-time feedback. Based on the findings of our initial studies and on the analysis of several reports providing information about usage practices of mobile devices, we combined that knowledge with the capabilities of current mobile platforms in sensing the environment. The outcome of our analysis has informed the list of contextual information supported in Buddy Tracker, which is presented in Table 1.

Buddy Tracker uses several available sensors, such as GPS, accelerometer, light sensor, Android system logs, information about currently running applications and other methods to collect the most accurate information about the user's context. Calendar entries can be used to determine the user's current activity; and Google Geo Service is used to translate GPS coordinates into more meaningful text descriptions. We use Skyhook's Core Engine SDK to collect information about the user's current position. Unfortunately, due to technical problems and a negative effect on battery life, we could not implement noise level sensing in our field trials, although this feature was implemented in our lab-based standalone application prototype.

Contextual information describing the user's current situation is used in the learning module, which is an integral part of the real-time feedback manager, and its role is to learn new rules based on the user's context.

3.4 Learning Engine and Rules Enforcement

At the heart of Buddy Tracker architecture is the learning engine. The learning system is responsible for analyzing the user's context and creating a predictive model of their preferences for the real-time feedback. We initially developed our own learning algorithm, which required users to provide additional information about the relevant context. In other words, in addition to the questions presented in Figure 5, participants also had to tell the system what contextual factors had an impact on their decision that the feedback notification was or was not appropriate to and satisfactory in their particular situation when they were notified. Preliminary tests showed that the algorithm was not usable because it required too much input from users. As a result people were less willing to answer questionnaires. Therefore we decided to find more affordable learning algorithm, which resulted in a more usable method of collecting users' feedback.

We evaluated several algorithms available in Weka data mining software [25] and we decided to use J48 implementation of C4.5 algorithm [26]. The algorithm is capable of learning from incomplete contextual information, which minimizes the user's effort and increases the chance of collecting useful data. We also found that this algorithm has an intuitive way of representing the rules, which is useful in examining the most relevant contextual factors. By most relevant context we mean elements of context that have the biggest impact on users' preferences. Other algorithms we considered performed poorly on a few preliminary tests and C4.5 was preferable because it is an established (and reliable) technique.

The learning engine consists of two main elements, the learner and the interpreter. The learner's job is to analyze users' feedback and cross reference it with contextual factors describing the environment and create a user's model. A model is a decision tree

that contains users' rules for real-time feedback preferences. The second element of the learning system, the Interpreter is responsible for rules enforcement. The interpreter can be seen as a function $f(C_1, C_2, \dots, C_n)$, where each argument describes individual contextual factors as shown in Table 1 (i.e. $C_1 = \text{location label}$, $C_2 = \text{ringer mode}$, $C_3 = \text{mobile activity}$ etc.). Complete function can be represented as $f(C_1, C_2, \dots, C_n) = \text{RTF}$, where RTF , the output of this function, is the most appropriate real-time feedback representation for the given situation.

4. CONTEXT-AWARE REAL-TIME FEEDBACK IN ACTION

We conducted a study in two parts aimed at exploring the role of context in minimizing the intrusiveness and increasing the effectiveness of the real-time feedback:

- An online survey, in which we collected information about users' feedback preferences based on their responses to potential scenarios.
- A field trial of Buddy Tracker, equipped with context-sensing capabilities and machine learning support for better user experience and accuracy of the feedback representation.

The first study helped us collect data, which we used to derive initial rules for users' preferences. We used those rules in the first phase of the field trial to minimize the intrusiveness of the technology and avoid potential risk of losing participants before we could collect information about their personal preferences.

4.1 Towards Context-Awareness in the Real-Time Feedback Manager – the Survey

We conducted a scenario-based online survey to collect information about people's preferences for the real-time feedback notifications. We were aware that using surveys to analyze users' experience isolated from their natural setting might introduce biases due to the discrepancy between peoples' beliefs and intentions, and their actual behavior [27]. Therefore, we used videos to allow users to indirectly experience the interface as realistically as possible. In this section we present the design, method and recruitment process. We also discuss how the findings of the survey informed the field trial of the real-time feedback technology.

4.1.1 Survey Design and Method

The main goal of our survey was to capture people's preferences for the real-time feedback notifications if they were using our Buddy Tracker technology. Based on the data we collected during interviews with users of location sharing technologies, focus group participants and information about people's mobile practices, we created 24 different scenarios. Each scenario was especially designed to cover different types of contextual information, such as location, real-world activity, mobile task, time of the day or presence of other people. For example, a scenario looking at users' feedback preferences in context would read: "Imagine, you are seated in a restaurant and, while waiting for your meal to arrive, check your bank statement via mobile browser. When, Jenny checks your location." In this scenario, the context comprises the user's location, phone's position, mobile activity and real-world activity. Results presenting users' preferences for this scenario are shown in Figure 4. In addition to the textual description, we also showed the participants an image illustrating the situation. A list of all scenarios is available at <http://bit.ly/RTFLocationSharingScenarios>.

Each participant was asked to express his or her preference for each of 10 scenarios, randomly assigned to them amongst the 24 available. Users had to assign at least one feedback representation to one of the three options: best choice, acceptable or "please no". Users could not assign one feedback type to more than one option, i.e. sound cannot be best and acceptable choice at the same time. We used 3 groups of preference options consistently with the requirements of the learning algorithms we used to derive rules, which use elimination criteria in rule generation. This approach also helped us identify most unacceptable feedback representations, which were ignored in case of small differences in certainty of positive answers.

Because many people are not familiar with Android platform and we did not want to limit our participants only to Android users, we used short videos to present how each interface works on the mobile device. Before the user could proceed to the questions, he was asked to read information about each interface and watch a short video presenting its functionality. We also added a video preview option in the answers section so users could see the interface working while answering the question, just by clicking on its name on the list. A screenshot of the survey is presented in Figure 3.

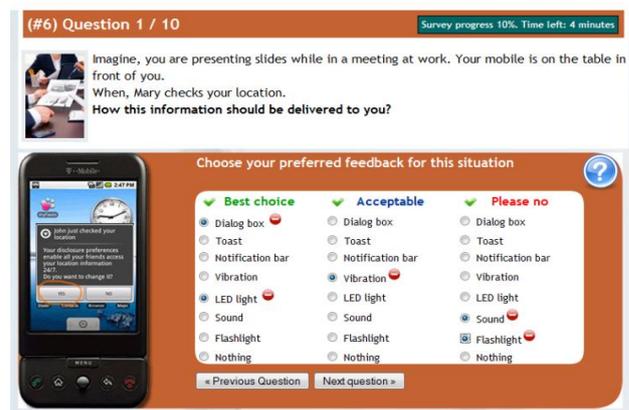


Figure 3. Screenshot of the survey presenting the question view. In this view the user is presented with the example scenario supported by the visual representation. Underneath, the user could see a mobile device, in which he can preview how the notification looks like in the real world. The last section allows the user to assign the notification of his choice to one of the three options.

4.1.2 Recruitment and Participants

We launched our survey in February 2010 and advertised our research through word of mouth, social networking sites (especially Facebook and Twitter) and location-sharing applications such as foursquare and Brightkite. We raffled a £45 (approx US\$70) gift voucher as compensation for taking part in the study.

A total of 3136 people started the survey. We discarded the data of all respondents that completed the survey in less than 4 minutes or did not provide answers to 10 scenarios. Due to the number of questions and instructions it was impossible to complete the study in less than 4 minutes, unless the respondent was already familiar with the interfaces and did not need any training. The following findings are from the 216 surveys that remained which give 2160 users' preferences for the real-time feedback in different situations.

The gender ratio of our respondents was 3-1 male to female. The age demographic showed that 46% (101) of all respondents were in their twenties, 29% (63) were aged between 31 and 40, 21% (45) over 41, and 3% (7) under 20 years old. Over 60% of all participants were active users of location sharing applications, mainly foursquare, Brightkite, Gowalla and Google Latitude. Interestingly, some of the respondents mentioned Twitter as location-sharing service.

4.1.3 Findings- Initial Rules for the Field Trial

The main goal of this work was to derive initial rules, which would minimize the negative impact on users' experience in the field trial (described in the next section). We also wanted to minimize the number of participants that could potentially withdraw from the study due to the system delivering inappropriate feedback. We mapped each scenario into the context model (set of contextual variables describing the situation that is supported by our system). We then used statistical models to generalize users' preferences and assign the most appropriate representation for real-time feedback to the context. We encoded generalized preferences as rules in the form of a decision table, which was used in the field evaluation of the technology. We also used this data to find the most appropriate learning algorithm for the field trial.

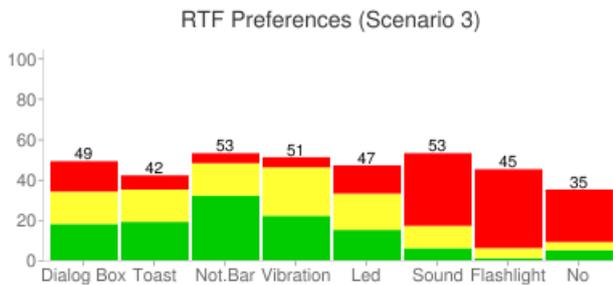


Figure 4. Survey results for scenario 3 (User at the restaurant received a notification). This chart shows that sound is not an appropriate in the situation; however experience from the field trial suggests that people are likely to accept more intrusive notifications for their entertainment value.

4.2 Field Trial

Initial rules describing users' preferences for the real-time feedback in a specific context were derived from the survey and we were ready to test our technology in the real-world. We performed a field trial aimed at evaluating the effectiveness of context-awareness and machine learning at improving the usability of the real-time feedback. Our objective here was to see if learning has an impact on users' experience and acceptance of the technology. We were also interested in assessing whether associating contextual factors with users' preferences can minimize the intrusiveness while maximizing the effectiveness of real-time notifications. Intrusiveness was the biggest criticism to this technology reported by participants of our previous studies [8].

4.2.1 Participants and devices

We recruited 15 participants (3 males and 12 females) from a variety of backgrounds and included a truck driver, a dental surgery assistant, a PhD student, a logistic officer, a sales manager, a graphic designer, curriculum managers, software developers, a CEO of a big company, a flight lieutenant from the Royal Air Force and a currently non-working person. Participants

could invite their friends to take part in the study as 'requesting' users. These used our web or iPhone app to check on their friends' location, but did not provide us any data about the real-time feedback, as their location was not being tracked or requested. Therefore we do not consider requesting users participants, as they were only requesting location of their friends in order to increase the number of realistic real-time notifications presented to participants actively sharing their location. Participants and their requesting users were split into seven groups. Most of the participants were based in the UK, one participant was based in Cyprus, one in Poland and one visited Vietnam during the study.

Setup

Prior to the study participants were asked to complete the Westin Harris privacy survey [28]. We also asked our participants to specify a set of meaningful locations, i.e. home or office, using the location manager module developed for the purpose of this study. Each user could create their own database of locations, assign them to one of the existing categories [24] or specify their own label for place. Users could also create areas by specifying a radius (between 50 yards and 2 miles). This information was used to correlate users' preferences with places and improve the learning process. Participants were instructed about additional settings provided by the system, such as quiet hours, that allowed them to create time periods when they did not want to be disturbed by the system. Users could also specify what types of notifications could be used. We also asked our participants to check if their mobile device met the requirements of the system, such as checking if TTS (Text To Speech) functionality was installed on their phone.

Four participants did not have their own Android device but they accepted devices from us, and used them as their main phones for the period of this study.

Participants were informed about the purpose of the study and information collected. We also made them aware of the negative implications of participating, such as reduced battery life. Extended batteries were offered for those participants that were not satisfied with the battery life on their phone. We also explained that we had instrumented the interface to collect information about any tracking events. Participants were offered £40 (approx \$65) for completing the 3-week study including post-study interviews, each lasting 30-90 minutes. Our research protocol was approved by our institution's Human Research Ethics Committee (equivalent to IRB).

4.2.2 Method

The study spanned a total of 3 weeks and was split into two phases. During both phases users received a real-time feedback notification every time someone looked-up their location; 5 minutes later they received a text message with a unique URL to the extended experience sampling questionnaire (see Figure 5). Users were asked to (1) score the accuracy of location where the system logged them at the time of sending the notification; (2) say if they noticed the notification; and (3) tell us if the notification used was appropriate for the given situation. Similarly to the survey presented in section 4.1, users had three options to score the notifications: YES (best choice), OK (acceptable) and NO (unacceptable notification). Users were asked to suggest a more appropriate notification for the given situation when they answered OK or NO. A drop down list of available types of feedback appeared, from which users could choose a better option. Information from questionnaires was used by the learning

engine to improve the accuracy of later notifications. Users could also specify a memory phrase, a unique textual description to help them recollect the specific situation in detail during debriefing session [29].

Figure 5 consists of two panels, A and B, showing a feedback form. Panel A is a compact version with three main sections: 'Situation', 'Real-time notification', and 'Memory Phrase'. Panel B is an expanded view of the 'Toast' notification section, showing a list of notification types for selection.

Panel A: Feedback Form

- Situation:** Real-time notification was sent to you 6 days ago because [redacted] looked-up your location. We thought that at the time of sending the request you've been at home sweet home [View on map]. Please tell us if this location is correct. (YES/NO buttons)
- Real-time notification:** You were notified about your location request via Toast. Have you noticed it? (YES/NO buttons). Was the Toast appropriate for the current situation? (YES/OK/NO buttons).
- Memory Phrase:** please enter a descriptor of the situation, a short message/tag that reminds you about the notification. (reading news, at home friday night)
- Submit my feedback** button.

Panel B: Expanded Toast Feedback

- Toast:** Have you noticed it? (YES/NO buttons).
- Was the Toast appropriate for the current situation?** (YES/OK/NO buttons).
- Not satisfied with the type of real-time notification provided? Please choose the one that works better for you:**
 - No feedback
 - Notification bar
 - Toast
 - dialog box
 - flashing LED light
 - Alert box
 - Vibration
 - Sound
 - Audio message in natural language

Figure 5. (a) Feedback form used to collect the data from users. Form consisted of two main sections: situation and real-time notification. Users could also specify a memory-phrase, a unique description that could remind him about the situation [29]. (b) Expanded view, presents additional drop down list of feedback types, which enabled users to instruct the system what would be more appropriate feedback representation in the given context.

During phase 1 we collected data about individual's preferences in order to create a training dataset for the learning engine. This phase lasted 2 weeks. In the second phase, lasting 1 week, we introduced real-time feedback notifications that were based on output from the learning engine. This was the only difference between the two phases.

At the end of the field trial, participants individually took part in post-study debriefing sessions. Interviews aimed at exploring the participants' decision making process, i.e. *why a particular type of notification was deemed to be preferable over another; or why participants expressed different preferences even when their context appeared to be the same (from the system's point of view)*. Materials used during debriefing sessions included (a) users' answers to experience sampling questionnaires; (b) system logs containing activities users' had undertaken during the study; and (c) results of the initial analysis performed on users' data, such as situations when people noticed a notification but scored it

negatively, or situations when people answered differently in the same context.

4.2.3 Findings

Over the period of 3 weeks the system sent 3937 experience sampling questionnaires (same as the number of real-time feedback notifications), of which 2192 were answered successfully. Participants started but not completed 114 questionnaires and 809 were ignored by participants. The overall participation rate, calculated as the percentage of successfully completed questionnaires, was 56%. Participants answered an average of 146.6 questionnaires (median=137). The most active participant completed 257 (user 7) questionnaires and the least active only 19 (user 32).

In this section we describe the main findings of our study. We use both quantitative and qualitative data collected during the study to report our findings. We begin by evaluating the efficiency of the context-awareness and learning engine at providing an unobtrusive and effective notification mechanism. Then we examine the effect of selected contextual information on the system's performance and users' acceptance of real-time feedback technology. We also discuss social issues related to real-time feedback, such as how the users' job affected the accuracy of notifications and discuss issues related to mobile privacy. We also look at how the group's privacy can be violated by our technology and how people overcome those problems in their life.

Acceptance of the Context-Aware Feedback

After examining the users' answers we noticed that 13 participants experienced an increase in the system's accuracy in the second phase of the study. By accuracy we mean the appropriateness of the notification used for the given situation. Only two users' reported lower accuracy in the second phase. User 12 and user 14 experienced 5% and 51% drop in the accuracy of notifications respectively (see Figure 6). We found that the system could not adapt to user 12 as he kept changing his rules due to his unpredictable circumstances and social context (he was a sales person travelling across the country). We describe the impact of lifestyle on the system's performance later in the paper. Additionally, the low accuracy of user 14 is the consequence of unexpected problems he experienced in the last phase; he could only participate for one day during the final week. For this reason we have omitted user 14 when calculating the average accuracy of the system.

The average accuracy of feedback delivery was 72.45% in the first phase, and 86.75% in the second phase. An average increase in the system's performance was 14%, the maximum increase was 46%. Surprisingly the user with the higher increase (user 21) did not notice a significant change in the system's performance.

The Impact of Learning on User Experience

Only a few users reported that they actually noticed an improvement in the system's accuracy. Some users were able to precisely point out the moment when the system started to provide more accurate feedback notifications. We found that better accuracy contributed towards greater reliability on the part of the system and trust on the part of the user.

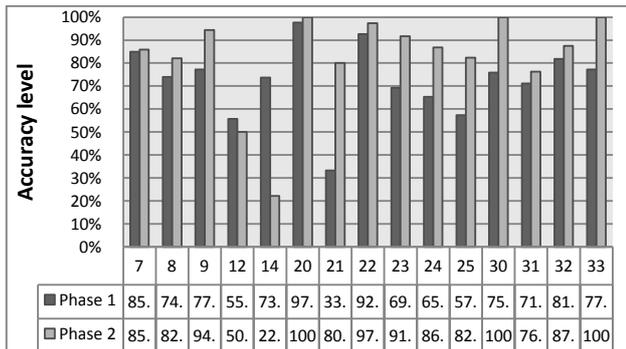


Figure 6. Accuracy of notifications during two phases of the study. Chart shows a percentage of positive feedback given by each user during 2 phases of the study.

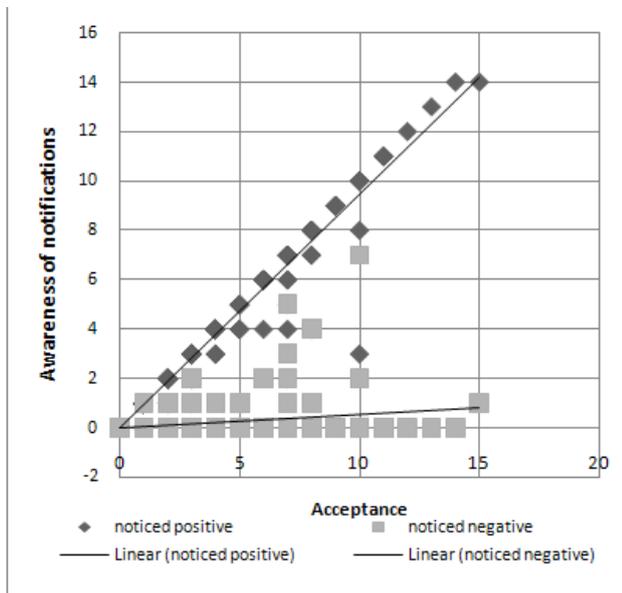


Figure 7. Correlation between the number of noticed notifications (awareness of notifications); and positive (diamond markers) and negative feedback (square markers). This chart suggests that acceptance of technology is correlated with awareness of notifications (correlation rate = 0.9788).

During post-study interviews, users reported that they felt more confident and comfortable using the system in the second phase. We found that awareness of the notification was a contributing factor towards the acceptance of this technology. In other words, the more notifications participants noticed, the more positive their feedback was. Figure 7 illustrates the correlation between the awareness rate and acceptance rate.

We observed 49 cases when people noticed a notification but scored it negatively, 42 in the first phase and 7 in the second phase. Of these, 37 were caused by more intrusive notifications (natural language, sound and vibration). Of these 33 were in the first phase and 3 in the second phase. This shows that the learning system helped us minimize the intrusiveness of this technology. We observed an 83% drop in the number of intrusive notifications presented to users in phase 2. During post-study interviews participants reported that most of these negative answers were the consequence of unpleasant situations caused by the system, for

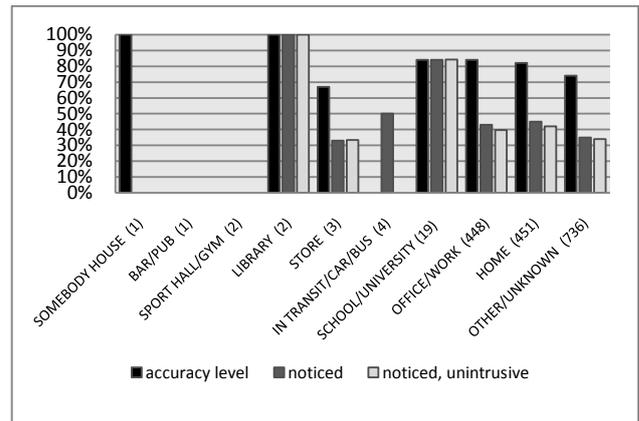


Figure 8. System performance vs. location. This chart presents the accuracy of notifications at different locations (black bar), percentage of noticed notifications (dark grey bar) and level of unobtrusiveness (light grey bar). The number of situations when users received a real-time feedback at the location is shown in parentheses.

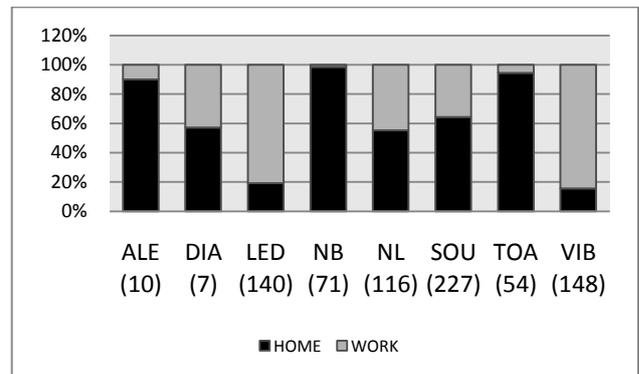


Figure 9. Users' preferences for real-time notifications at home and work using different methods (abbreviated). This chart shows the percentage of positively scored notifications for two most frequent known locations with the total number of situations for each method shown in parentheses.

example, if natural language notification was used during a meeting or while having dinner with friends.

The awareness rate was increased by 7.29% in the second phase (46.98%). By awareness rate we mean the percentage of noticed notifications. This means that the learning system helped us to not only increase the unobtrusiveness but also increase the efficiency of information delivery, which had a positive impact on usability and users' acceptance of the technology. The increased awareness contributed towards greater trustworthiness and reliability of the system.

Effect of Location on Acceptance and Awareness

Figure 8 shows the distribution of users' locations at the time when they received real-time feedback notifications. Most of the time participants received notifications while they were at home, at work or in an unknown place. The chart suggests that there were only a few situations when participants were notified about look-up events while at a store, pub or in the library. However, many of these locations are covered in the other/unknown group. Prior to the study, participants found it difficult to specify all the

places they were likely to visit during the study, therefore we cannot perform a detailed analysis of users' preferences for the real-time feedback in relation to the contextual factor location. Instead, to explore the effect of location on users' preferences, we use the two most frequent known locations, namely home and work.

The usage rate (number of situations when users actively used their phone at the time a notification was received) was higher at home, which suggests why toast (TOA) and notification bar (NB) visualizations were preferred at home. We observed that for 47% of notifications delivered while users were at home, the participant was using their phone at the time. The usage rate at work was 32%. The most commonly used feedback representations at work were LED and vibration (VIB). Surprisingly, we did not observe a big difference in the acceptance rate of natural language notifications (NL) between the two locations (See Figure 9). An interesting observation is that 95% of accepted NL notifications at work were reported by office workers with more than 10 people in the office. However, only three positive NL notifications, when at work, were reported by the participant who was a truck driver.

Effect of Mobile Activity on Acceptance

We noted that our participants used a range of 24 different applications, at the time at which a real-time feedback notification was received. Mostly these were social and communication applications. We examined the decision tree models generated for each user by the learning algorithm and found that current mobile activity was the common node for 10 users. Our analysis shows that mobile activity is an important element of context, which can be used to determine the most appropriate type of notification while the phone is in use. This is especially important for the system designers, as it can minimize the contextual information that must be collected through battery-consuming sensors.

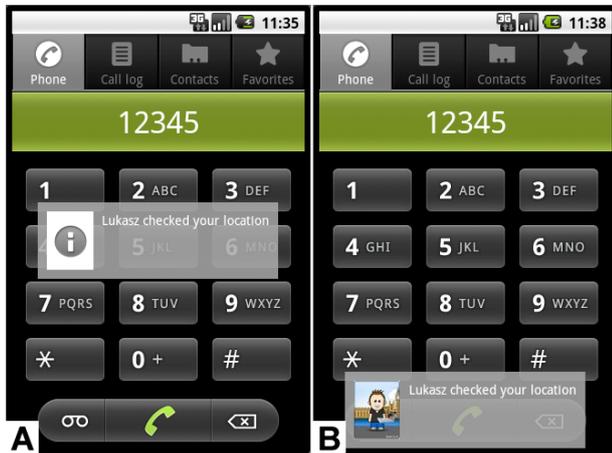


Figure 10. Toast notification, a small, semitransparent floating window appearing on the screen and disappearing after 2 seconds. Pictures presents toast displayed to the user while typing a number. (a) Less usable and more intrusive, toast in the first phase placed in the middle of the screen; (b) more usable, toast in the second phase displayed in the bottom of the screen, a picture of requester was presented in the second phase.

One interesting observation was that even less intrusive notifications, such as toast (see Figure 10), are perceived as

annoying while watching videos, browsing the Internet or typing. We observed that even people with a pragmatic approach to the technology did not like to be disturbed when performing one of the above tasks.

In the second phase we have changed the position of the toast notification, and aligned it to the bottom of the screen to see if the main factor was really the current task, or if we could solve this problem by changing the position of the notification (see Figure 10). In the second phase we also displayed a picture of the requester to minimize the cognitive effort of assimilating the information. We observed that in the second phase that bottom alignment had a positive effect on user experience while browsing the Internet; no negative feedback for toast notifications was reported in the second phase. Users were not presented toast notifications while watching videos or typing in the second phase. However data collected during interviews confirms that alignment of the notification and the new design has a positive effect on user experience, but there are situations in which the system should respect people's privacy and not display any visual notification that covers the working area of the screen (i.e. while people type text or watch videos).

Lifestyle and Personality vs. System's Performance

Our study suggests that real-time feedback technology is not for everyone. Some people are just not interested in who viewed their location. Others do not seem to like to know this information in real time (however, this attitude might change based on who's checking and how important it is). Finally, others report more practical reasons for their low level of acceptance, such as battery consumption. We looked at the acceptance rate in relation to users' occupation and other life patterns. We found that it is more difficult to provide both an unobtrusive and effective notification system for mobile occupations (for example in cases such as those of the truck driver or sales person). We observed that a proximity sensor showing the distance between the owner and the mobile device would solve many performance problems in the case of highly mobile users.

We also found that system acceptance depends on two main factors: intrusiveness of the method of delivery and effectiveness (awareness of notifications). Although effectiveness is a positive factor contributing to the acceptance of real-time feedback technology, nevertheless it is not important for everyone. We found that effectiveness is more important for people with a pragmatic approach towards technology while intrusiveness is more important for people with a more opportunistic attitude. There is a stronger need for historical feedback in the second group than there is in the first group.

Attitudes towards location sharing technologies

Amongst the 15 participants taking part in our study only 4 had prior experience with location sharing and tracking technologies. When asked about their attitude towards the Buddy Tracker technology in the debriefing interviews, participants reported different reasons for sharing, or not, their location and for tracking others.

The most common reasons for using the Buddy Tracker was to support coordination; ensure a loved one's safety, e.g., tracking children, monitoring a partner's long journey; or just curiosity. Additionally, some users reported using Buddy Tracker and the real-time feedback technology as a game playing activity: for example, one of the participants reported that "when someone

looks me up, I look him up in return". Some people used the system to check where they have been or track their running activities. One of our participants mentioned that very often he shares his tracks with other people having interest in sport. He called that *conversational sharing*, during which he discusses his achievements using location data. Despite the socially-oriented aspects of using location sharing technologies, we found that interest in location sharing might have a non-social background. We noticed that one of our participants used the technology because he likes maps. A real-time map view allowed him to monitor how his buddies travelled between their work and home destinations. People also used Buddy Tracker to coordinate their life, i.e. check if the person is "contactable" (for example, one of our participants was in Vietnam during the study, and sometimes had problems with connectivity) or express feelings (for example, a female participant tracked her partner via Buddy Tracker to coordinate her cooking with her partner's journey home in order impress him with a hot dinner upon his arrival).

The most common reasons for not using location sharing technologies are lack of interest in someone else's life, ethical issues related to tracking and what appeared to be the less socially-orientated personality of some participants. People also described several technical problems that had an impact on their acceptance of this technology, such as battery life and accuracy of location, both of which make the system less useful. Those who reported negative aspects of location sharing were mainly people with a utilitarian and pragmatic approach to location sharing.

Real-Time Feedback and Mobile Privacy

We designed the real-time feedback technology as the tool for supporting awareness, which can help people understand the actual data flow and help them make more informed privacy decisions. What we mean by personal privacy in this context is *the processes by which individuals selectively disclose personal information* [18]. However there is much more to privacy than management and rules setting, and this study sheds a new light on aspects of privacy specific to mobility and mobile technology. We revealed several mobility issues that have an impact on our personal privacy and sense of solitude.

Mobile phones are not communication-only devices anymore; our participants described their devices as "a *computer that happens to be a phone*" or an "*information device*". Some of them used their phone only for communication, but most of our participants used their phones for many different purposes, such as entertainment, listening to music, personal organization, networking, checking emails, news reading or physical and virtual navigation. Most of our participants were in the close proximity of their phones 24 hours a day. Only 4 participants reported that they keep their phone in a bag, in the kitchen or in the living room over night. Understanding the context in which people use their devices, seems to be very important when analyzing the privacy impact of real-time feedback technologies, and any notification mechanisms in general.

Buddy Tracker used several representations for real-time feedback to minimize the negative effect of notifications such as intrusiveness. However there were situations where even less intrusive notifications, such as toast, had a negative impact on a users' experience and their sense of privacy. During the debriefing session we asked one of the participant (female, 28) why she refused the toast notification while she was watching a video on YouTube (prior to that episode she rated this type of notification

positively 43 times). She reported that at that moment she was relaxing with her 2 years old son and they were watching television on her phone. The notification was not only annoying but also caused a bit of anger and also affected her sense of privacy as she did not want anything or anyone to disturb her. She tried to dismiss it immediately but she had no control over it, and had to wait 2 seconds before the notification disappeared automatically. When asked about this experience, the user reported that "*there are times when a phone should not disturb and should not act as an information device*". A similar incident took place with another user (a 28 year-old male student), while he was preparing for his supervision meeting on the next day. He was notified about a location look-up event via audio message in natural language. He remembered this situation as one of the most annoying during the whole study as at that moment he was in a bad mood and needed to withdraw from his surroundings. The intrusion to one's privacy caused by any notifications, by real-time location feedback, incoming SMS buzz or a ringtone, can have a negative effect and could cause frustration or even embarrassment, all of which has an impact on our willingness to accept a technology. On one hand, technology is useful and desired, on the other hand, it can disturb and affect one's *right be let alone* [30].

A key observation here is that, when we talk about mobile privacy, we should distinguish between the control layer, supporting the privacy of information generated or shared through the mobile service, and the communication layer, between the service and the user. The latter, responsible for information presentation, can have an impact on different aspects of privacy, for example, intimacy and solitude as the incidents with our participants indicate. Mobile privacy is not always about protecting information, but it also refers to a mobile service user's state of mind, which determines when and how to interact with the environment. It is also important that in these instances privacy was being violated even if the phone was not being used. In other words, privacy-sensitive applications need to take into consideration the *entire* context of the user. However, given the technological limitations of current mobile platforms, this can only be approximated at this time.

Real-Time Feedback and External Image

The post-study interviews revealed that data owners, those that receive a notification when somebody looks-up their location, were concerned about their external image while using the technology. For example, inappropriate notifications presented during a meeting may be regarded as a sign of disrespect to other people and the device owner could be perceived as rude. Therefore the criteria for successful notification include the user's imagined view of how others around them will react. In other words, this means that people choose notifications that would be good for them but also acceptable to people nearby. Participants reported a number of unpleasant situations when an inappropriate notification was used, for example, when natural language was used during a meeting, which irritated others and embarrassed the user. Disturbing other people was a common example of the technology's intrusiveness. However, 5 people, those with more utilitarian view on location sharing and real-time feedback technology, reported that they would accept more intrusive notification for improved awareness.

Our study shows that people make judgments about others based on how they use technology. One user reported that "*someone, whose phone is talking all the time does not make a good*

impression and seems to be unreliable person”, which suggests that delivering notifications too frequently might have an impact on the user’s external image.

We asked participants to use all available types of notifications and let the phone learn from their behavior what the most appropriate notification was in any given context. However, to provide additional control over real-time feedback and minimize the intrusiveness, Buddy Tracker allowed participants to limit the types of notifications which could be used. We observed that some people switched off the most intrusive notifications (vibration, sound and natural language), some of whom did it to protect their reputation by not allowing the technology to do things that are not acceptable in the given environment, e.g., the work place. Many users kept their phone in silent mode, which automatically disabled the audible feedback in the Buddy Tracker.

5. DISCUSSION

We have presented our work on context-aware real-time feedback, a novel application for supporting awareness in privacy-sensitive mobile applications. Our main objective was to design and develop an unobtrusive notification mechanism empowering users by providing them with information about their location data flow. The real-time feedback supports them in the ongoing process of privacy management and thus helps them to enjoy the social participation afforded by ubiquitous technology with awareness and proactivity. We used rich contextual information and machine learning techniques to adapt the system to each individual in order to improve the acceptance of the technology and minimize the intrusiveness of notifications. The findings of our study show that our approach can significantly minimize the intrusiveness, which has a positive impact on user experience.

They also suggest that designing unobtrusive notifications for ubiquitous technology is a very challenging task due to several factors, such as technological limitations (battery problems, accelerometer not working while phone in sleep mode, lack of proximity sensor), personal attitudes, occupation type or mental state (concentration, mood).

We decided to use off the shelf mobile devices in order to increase the realism of the study. Although modern devices are capable of sensing rich information about the users’ environment, battery consumption of sensors make it difficult to develop a long-lasting and reliable system. Due to this problem we could not use all available sensors to provide richer contextual clues to the learning system, e.g., microphone could not be used to measure the noise level. Surprisingly, our participants reported technical limitations and high battery consumption to be the most negative factors towards the acceptance of the technology. This indicates that that current technology is not yet able to provide *usable context-awareness*.

To the best of our knowledge, our study was the first to explore the use of context-aware notification mechanisms supporting awareness in the field of mobile computing. Our findings indicate that contextual factors are critical in determining a user’s reaction to different forms of feedback and that the ability of a device to learn about and adapt to a user’s context, afforded by machine learning, can determine the acceptance of real-time feedback technology. Although the study was small and most of our participants were males, our findings provide new insights about the actual impact of real-time feedback technology on mobile privacy and mobile privacy management.

However, our findings also suggest that there is much more to privacy in ubiquitous computing than control and management of data flow. We observed an interesting phenomenon about mobile privacy within the spectrum of human computer interaction. Users’ acceptance of the technology depends not only on the intrusiveness and effectiveness of notifications from a pragmatic or social perspective, but also on what we might call the users’ *emotional context*, which might include for example their level of concentration, or their need for intimacy or solitude, at notification time. These factors have a significant impact on how users perceive technology, and whether and how they want to be alerted.

The application presented in this paper uses a client-server architecture, which is not ideal due to potential security problems (i.e. a third party could potentially monitor the traffic to collect data transmitted between the client and the server). Technological improvements of Buddy Tracker’s architecture are needed and should include using a client’s device for data capture, data storage and learning process in order to minimize privacy and security problems related to data transfer. One of the next items in our research agenda is to use on-device learning approaches, similar to Wang [31].

Despite the positive findings of our study, we need to investigate further *how to incorporate machine learning into this type of systems in a usable manner*. As reported here, we managed to improve the accuracy of notifications and significantly minimize the intrusiveness of the technology. However, similarly to work by Sadeh et al. [20], our methodology required users’ feedback to learn new rules, which might be cumbersome. Although we were able to use an algorithm that is capable of learning rules from a limited data set (hence minimizing user effort), incorporating the learning process into real systems still poses a significant challenge. While people may be willing to give feedback to the system while taking part in an experiment for which they receive compensation, they may not be willing to do the same in an uncontrolled setting. Therefore the next step in this research is to explore any links between users’ personalities, occupations, and life style in general, and privacy attitudes in order to determine whether an adaptive model for user acceptance of real-time feedback technology can be derived. Concomitantly, new input methods may be required, which will allow us to design more affordable ways of defining learning rules.

6. CONCLUSIONS

We have conducted a study investigating the potential of using contextual cues and machine learning to increase the usability of real-time feedback technology. Based on both quantitative and qualitative data, collected during a 3-week field trial and following debriefing interviews, we reported several novel findings. We showed that machine learning and context-awareness can increase the contextual appropriateness of notifications, which contributes towards greater trust in the system and higher level of comfort, thereby increasing the overall user experience. While we observed a significant drop in the intrusiveness of the system in the second phase, we did not observe a significant increase in the effectiveness of notifications.

We also observed a number of social implications related to real-time feedback technology, such as impact on mobile privacy and social image. Our findings indicate that mobile activity has an impact on users’ preferences, and information about users’ mobile activity is useful in determining an appropriate notification type

for a given context. Although machine learning techniques are efficient at improving user experience, new methods for collecting users' feedback are needed to make the learning process more transparent and more usable in real-world applications.

7. ACKNOWLEDGMENTS

We are indebted to the participants who took part in the study for their invaluable contribution. This work is funded by the EPSRC grant EP/F024037/1.

8. REFERENCES

- [1] I. Altman, *The environment and social behavior*. Brooks/Cole Monterey, CA, 1975.
- [2] S. T. Margulis, "On the Status and Contribution of Westin's and Altman's Theories of Privacy," *Journal of Social Issues*, vol. 59, no. 2, pp. 411-429, 2003.
- [3] J. Archea, "The place of architectural factors in behavioral theories of privacy," *Journal of Social Issues*, vol. 33, no. 3, pp. 116-137, 1977.
- [4] V. W. Kupritz, "Privacy management at work: A conceptual model," *Journal of Architectural and Planning Research*, vol. 17, no. 1, pp. 47-63, 2000.
- [5] H. Heft, *Ecological psychology in context: James Gibson, Roger Barker, and the legacy of William James's radical empiricism*. Lawrence Erlbaum, 2001.
- [6] P. Dourish and V. Bellotti, "Awareness and coordination in shared workspaces," in *Proceedings of CSCW '92*, pp. 107-114, 1992.
- [7] T. Erickson and W. A. Kellogg, "Social translucence: an approach to designing systems that support social processes," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 7, no. 1, pp. 59-83, 2000.
- [8] L. Jędrzejczyk, B. A. Price, A. K. Bandara, and B. Nuseibeh, "On the impact of real-time feedback on users' behaviour in mobile location-sharing applications," in *Proceedings of SOUPS '10*, 2010.
- [9] J. Y. Tsai, P. Kelley, P. Drielsma, L. F. Cranor, J. Hong, and N. Sadeh, "Who's viewed you?: the impact of feedback in a mobile location-sharing application," in *Proceedings of CHI '09*, pp. 2003-2012, 2009.
- [10] L. Palen and P. Dourish, "Unpacking "privacy" for a networked world," in *In Proceedings of CHI '03*, pp. 129-136, 2003.
- [11] "Locaccino." [Online]. Available: <http://locaccino.org/>. [Accessed: 09-Feb-2011].
- [12] "Google Latitude." [Online]. Available: <https://www.google.com/latitude>. [Accessed: 09-Feb-2011].
- [13] M. Langheinrich, "A Privacy Awareness System for Ubiquitous Computing Environments," in *Proceedings of UbiComp '02*, pp. 237-245, 2002.
- [14] V. Bellotti and A. Sellen, "Design for privacy in ubiquitous computing environments," in *Proceedings of ECSCW '93*, pp. 77-92, 1993.
- [15] M. Raento and A. Oulasvirta, "Privacy management for social awareness applications," in *Proceedings of CAPS '05 Workshop*, 2005.
- [16] G. Hsieh, K. P. Tang, W. Y. Low, and J. I. Hong, "Field deployment of IMBuddy: a study of privacy control and feedback mechanisms for contextual IM," in *Proceedings of UbiComp '07*, pp. 91-108, 2007.
- [17] A. Sellen, R. Eardley, S. Izadi, and R. Harper, "The whereabouts clock: early testing of a situated awareness device," in *Proceedings of CHI '06*, p. 1312, 2006.
- [18] J. Hong, "An Architecture for Privacy-Sensitive Ubiquitous Computing," Unpublished PhD Thesis, University of California at Berkeley, Computer Science Division, 2005.
- [19] S. Lederer, I. Hong, K. Dey, and A. Landay, "Personal privacy through understanding and action: five pitfalls for designers," *The Journal of Personal and Ubiquitous Computing*, vol. 8, no. 6, pp. 440-454, 2004.
- [20] N. Sadeh et al., "Understanding and capturing people's privacy policies in a people finder application," *The Journal of Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 401-412, Aug. 2009.
- [21] D. H. Nguyen and E. D. Mynatt, *Privacy mirrors: understanding and shaping socio-technical ubiquitous computing systems*. GVU Technical Report: GIT-GVU-02-16: Georgia Institute of Technology, 2002.
- [22] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggle, "Towards a Better Understanding of Context and Context-Awareness," in *Proceedings on HUC '99*, pp. 304-307, 1999.
- [23] M. S. Ackerman and L. Cranor, "Privacy critics: UI components to safeguard users' privacy," in *Proceedings of CHI '99*, pp. 258-259, 1999.
- [24] A. Khalil and K. Connelly, "Context-aware telephony: privacy preferences and sharing patterns," in *Proceedings of CSCW '06*, pp. 469-478, 2006.
- [25] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The WEKA data mining software: An update," *ACM SIGKDD Explorations Newsletter*, vol. 11, no. 1, pp. 10-18, 2009.
- [26] J. R. Quinlan, *C4. 5: programs for machine learning*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc, 1993.
- [27] C. Jensen, C. Potts, and C. Jensen, "Privacy practices of Internet users: self-reports versus observed behavior," *International Journal of Human-Computer Studies*, vol. 63, no. 1, pp. 203-227, 2005.
- [28] H. Taylor, "Most people are "privacy pragmatists" who, while concerned about privacy, will sometimes trade it off for other benefits," *The Harris Poll*, vol. 17, p. 19, 2003.
- [29] C. Mancini et al., "From spaces to places: emerging contexts in mobile privacy," in *Proceedings of UbiComp '09*, pp. 1-10, 2009.
- [30] S. D. Warren and L. D. Brandeis, *The Right to Privacy*. *Harvard Law Review*. Boston: , 1890, pp. 193-220.
- [31] A. I. Wang and Q. K. Ahmad, "CAMF - Context-Aware Machine learning Framework for Android," in *Proceedings of SEA '10*, 2010.