

The Effect of Privacy Concerns on Privacy Recommenders

Yuchen Zhao

University of St Andrews, UK
yz39@st-andrews.ac.uk

Juan Ye

University of St Andrews, UK
jy31@st-andrews.ac.uk

Tristan Henderson

University of St Andrews, UK
tnhh@st-andrews.ac.uk

ABSTRACT

Location-sharing services such as Facebook and Foursquare/Swarm have become increasingly popular, due to the ease at which users can share their locations, and participate in services, games and other applications that leverage these locations. But it is important for people who use these services to configure appropriate location-privacy preferences so that they can control to whom they want to share their location information. Manually configuring these preferences may be burdensome and confusing, and so location-privacy preference recommenders based on crowdsourcing preferences from other users have been proposed. Whether people will accept the recommended preferences acquired from other users, who they may not know or trust, has not, however, been investigated.

In this paper, we present a user experiment ($n=99$) to explore what factors influence people's acceptance of location-privacy preference recommenders. We find that 44% of our participants have privacy concerns about such recommenders. These concerns are shown to have a negative effect ($p < 0.001$) on their acceptance of the recommendations and their satisfaction about their choices. Furthermore, users' acceptance of recommenders varies according to both context and recommendations being made. Our findings are potentially useful to designers of location-sharing services and privacy recommenders.

ACM Classification Keywords

K.4.1 Computers and Society: Public Policy Issues—*Privacy*

Author Keywords

location-based services; location-sharing services; privacy preferences; recommender systems; user acceptance

INTRODUCTION

With the help of social networking sites, location-based services, and smartphones, location-sharing services (LSSs), such as Foursquare/Swarm and Facebook's location check-in mechanism, have been evolving rapidly in recent years. People can now easily share their whereabouts with others, including their families, friends, and even the public, intentionally or unexpectedly. On the one hand, LSSs provide us with convenient features such as personal tracking and discovery

of new friends [19]. On the other hand, location-privacy preferences are highly context-aware [1]; sharing information depends on *where* they are and *when* they are in a certain place. Such context-awareness makes the self-configuration of privacy preferences a complex task [9], which may lead to a failure to protect location privacy [20]. Since people treat their location as the most valuable type of personal information [26], the lack of usable location-privacy protection mechanisms may affect the adoption of LSSs. To address these usability and information overload issues, recommenders have been designed to configure location-privacy preferences (semi-)automatically, for instance by recommending preferences based on crowdsourced results (e.g. from people with similar preferences). But people's acceptance of these recommenders and the factors that influence their acceptance have not been investigated sufficiently.

Existing research has mainly focussed on the measurement of metrics, such as the recommendation accuracy, by conducting offline recommender evaluation (i.e. using previously collected user data to construct and test a recommender). But location-privacy preferences are a relatively sensitive domain, compared with other fields where recommenders are used, such as music and movie recommendations, and so poor recommendations could have more serious consequences. Thus simply measuring accuracy may be insufficient. We also need to consider whether people accept these recommenders, or indeed the recommendations made by them. Such acceptance might be affected by subjective factors such as a person's concerns about sharing their preference data with a recommender, and as well as objective factors such as how the recommendations are made, or the type of recommendation being made. Therefore, it is necessary to investigate the influence of both people's subjective factors and these objective factors concerned with the recommender. Understanding this will help us decide which factors we should take into account when designing our recommenders.

In this paper we investigate the following research questions:

- **Q1.** Which subjective factors (e.g. trust in technology, privacy concerns, and perceived quality) will influence a person's acceptance of location-privacy preference recommendations?
- **Q2.** Which objective factors of recommendations (e.g. context and openness) will influence people's acceptance of location-privacy preference recommendations?

The contributions of our work are as follows:

- Using an online user study ($n=99$), we find that people do indeed have concerns about using recommenders for

privacy-preference prediction. These concerns have a negative influence on their perceived recommendation quality, satisfaction about their choices, and acceptance of the recommendations.

- The openness of a recommendation (with whom the recommender is proposing to share location) has a significant influence on people's acceptance. Recommendations with the highest and lowest openness are both less likely to be accepted.
- Context (specifically time and location) has an effect on people's acceptance of location-privacy preference recommendations.

We believe that our findings are of use to system designers who wish to use recommenders for configuring preferences about sensitive data.

RELATED WORK

To address usability issues in privacy protection, various machine learning classifiers have been used to predict privacy preferences in online social networks [8, 22, 7] and in LSSs [25, 2], thereby helping people to configure their privacy rules (semi-)automatically. These methods learn from individual users' previous privacy decisions and make predictions based on the models that they build. Experimental results demonstrate that these methods can achieve high prediction accuracy.

An alternative approach is to use crowdsourcing knowledge to make recommendations for privacy preferences, such as security configurations on mobile devices [12, 15] and location-sharing preferences [28, 29]. Compared with learning from individual histories, crowdsourcing recommenders learn from the "Wisdom of Crowds". For example, Toch's Super-Ego framework makes predictions based on semantic analysis of locations and the crowdsourcing results in the same semantic categories [27]. In addition, context-aware recommenders [3] based on collaborative filtering (CF) are also used, in which the location and time are treated as contexts and the recommendations are location-sharing preferences. For instance, Xie et al. combine both user-based and item-based CF in a recommender to make location-sharing preference recommendations [28]. They show that crowdsourcing methods can perform as well as machine learning classifiers. Furthermore, such CF recommenders can be used to bootstrap new systems, as they perform particularly well when there are insufficient training data [29].

Research has shown that people's security and privacy behaviours can be affected by social influence [6]. People often trust inaccurate recommendations more than they should [11]. Given that privacy-preference is a sensitive domain, it is important to investigate what factors would influence people's perception on the recommendations and their interactions with the recommenders. Knijnenburg and Jin [17] show that privacy recommendations for location sharing services have a strong persuasive effect on people. We extend their work by exploring what kind of factors from both the recommendations (e.g. source of crowdsourcing and level of openness

of recommendations) and the users (e.g. their concerns and perceived recommendation quality) can influence the users' acceptance of the recommendations.

Models and frameworks have been proposed to investigate the effects from other factors beyond recommendation accuracy (e.g. personal characteristics, perceived quality, transparency) in what is referred to as HRI (Human-Recommender Interaction) [21]. For example, Zins and Bauernfeind investigate which factors influence users' satisfaction with recommenders [30]. Their user study mainly focuses on effects from personal characteristics including Internet expertise, product involvement, and Internet purchase attitudes. By applying their model they find that these personal characteristics can influence people's experience when using online recommenders. Pu et al. propose a framework to evaluate recommenders from users' perspectives. They consider a variety of criteria including perceived system qualities, beliefs and attitudes, but do not take into account the influence from the recommender itself [23]. Compared with other existing frameworks, the user-centric evaluation framework proposed by Knijnenburg et al. [18] provides a structured model that covers all the factors in which we are interested, including both the objective factors of the recommender, as well as people's perception, experience and interaction, and personal and situational characteristics. We choose to use this framework, as described in the next section.

METHODOLOGY

Our goal is to investigate the effects of different factors on people's acceptance of location-privacy preference recommendations. To do this, we have designed an experiment based on the framework [18], which allows us to explore the relationships between various objective and subjective factors and the participants' acceptance.

User-centric evaluation of recommenders

In HRI, the evaluation of recommenders not only takes into account the objective factors (e.g. recommendation accuracy) but also the evaluation from the users' point of view (e.g. how good do the users think the recommender is). To study this, a valid framework is needed to precisely describe the factors of the evaluation by the users and build the relationships between these two parts. Knijnenburg et al. [18] have proposed a framework for such user-centric evaluation. The framework provides a way to measure the influences of the objective factors on the users' behaviours in HRI and these influences are considered to be moderated by users' subjective factors.

In this framework, a recommender is defined as a set of Objective System Aspects (OSA) that relate to the underlying recommendation algorithms and graphical user interfaces. OSAs influence the users' perception (e.g. perceived quality of the recommendations) of the recommender. To describe the users' subjective factors in the framework, their perception is defined through a set of Subjective System Aspects (SSA), which have influence on their Experience (EXP) (e.g. their satisfaction about their choices) and Interaction (INT) (e.g. purchasing the recommended products). The SSA is used as the moderators for the OSA's influence on EXP and

INT, which means the influence from OSA to EXP and INT is through SSA. The effects from the Situational Characteristics (SC) such as users' privacy concerns and from their Personal Characteristics (PC) such as demographics and domain knowledge to EXP and INT are also considered in this framework.

To apply this framework to location-privacy preference recommendations, we measure the following factors that might have an effect on acceptance of recommendations:

- *trust*: participants' general trust in technology.
- *quality*: participants' perceived quality of the recommended location-privacy preferences.
- *satisfaction*: participants' satisfaction about their chosen recommendations in the second part of experiment.
- *concern*: participants' privacy concern about using location-privacy preference recommenders.

Specifically we consider *trust* as PC, *quality* as SSA, *satisfaction* as EXP, and *concern* as SC. We use the different sources of crowdsourcing recommendations as conditions, i.e. OSA in the framework. We also introduce the participants' acceptance of the location-privacy preference recommendations into our model as an INT:

- *acceptance*: the percentage of the accepted location-privacy preference recommendations of a participant.

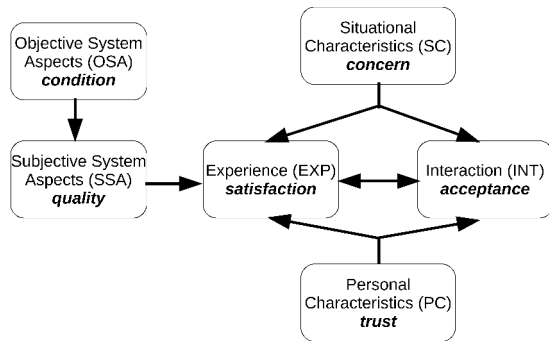


Figure 1. Diagram of the framework for the user-centric evaluation of recommenders as used in our experiment.

Figure 1 shows a diagram of how we apply this framework to privacy-preference recommenders.

Questionnaires

To evaluate the subjective factors in our experiment, we draw on the questionnaires provided by the framework [18]. To measure the above listed factors, we use the following questionnaires: *General trust in technology*, *Perceived recommendation quality*, *Choice satisfaction*, and *System-specific privacy concern* respectively. Each of these questionnaires contain several questions, each of which refers to a *question item*, measured on a five-point Likert scale from *Strongly disagree* to *Strongly agree*.

EXPERIMENTAL DESIGN

We hypothesise that people's acceptance of location-privacy preference recommendations change with different crowdsourcing sources, recommendations with different level of openness, and different contexts (i.e. location category and time). To test this, we conducted an experiment where people used three different location-privacy preference recommenders, and then studied which factors had an influence on their acceptance of the recommended location-privacy preferences.

Participants were invited to login with their Facebook accounts to our experimental system, so that our recommenders could use their real-world Facebook logins to generate recommendations. They were then provided with a prebriefing explaining the various recommenders, and how these recommenders could help with configuring privacy preferences. Next, they were presented with a series of recommendations about locations that existed in their Facebook location check-in histories (as shown in Figure 2). Each recommendation was presented as if it had been made using a particular source of data. For instance, in Figure 2 we see a recommendation that was generated using information from a participant's Facebook friends (in actual fact all recommendations were made randomly, so as to maintain objective accuracy without having it influence *quality*, *satisfaction*, or *acceptance*). Finally, participants were asked to complete a series of questionnaires about their perceived quality of the recommendations, their satisfaction, and system-specific concerns (as shown in Figure 3).

To recruit participants, we advertised our experiment through university mail lists and Facebook groups (to avoid biasing the samples, we used "location-sharing preference" instead of "location-privacy preference" in our experiment and advertisements). 164 participants tried to access our experiment and 99 of them had at least 10 distinct location check-in histories on their Facebook and completed the experiment. Each participant who completed the experiment received a £5 Amazon voucher for their participation. The study design was scrutinised and approved by our institutional ethics committee.

Prebriefing

Our prebriefing explained the various recommenders, which were:

- *same-location* recommender: using the preferences of people who have been to the same location;
- *similar-people* recommender: using the preferences of people who have similar previous location-sharing preferences;
- *Facebook-friends* recommender: using the preferences of people's Facebook friends.

Participants were provided with three examples to familiarise themselves with the recommenders, and a quiz was used to ensure that they understood the concepts. After successfully completing the quiz, participants were asked to login with their Facebook accounts using the PRISONER platform that

has been designed for privacy-sensitive social media experiments [14]. This ensured that we would only collect the minimum amount of data required for our purposes. To make our proposed recommenders look realistic, we asked for additional permissions such as a Facebook friends list, even though this was not needed.

To guarantee that we had enough data to generate recommendations, only participants with more than ten distinct location check-ins could proceed to the next part of the experiment.

Exploring recommendations

The purpose of the second part of our experiment was to measure the participants' acceptance of the recommended location-privacy preferences. Firstly, participants were asked about their demographic information including age and gender. Then we used a questionnaire to measure *trust*, one of the subjective factors in our experiment. The other questionnaires about *quality*, *satisfaction*, and *concern* were shown after the participants went through all the recommendations because we want to know whether these factors would be influenced after the participants using different recommenders.

Conditions were altered on a within-subjects basis. Each participant was presented with 30 location-privacy preference recommendations (10 for each recommender). For each recommender, we randomly selected 10 location check-ins from the participant's Facebook data as the contexts of the recommendations. For each recommendation (Figure 2), the participant was provided with the context of their previous check-in, a map of the location of this check-in, and the recommended preference from one of our recommenders. The context contained the name of the location and the time slot (i.e. morning, noon, afternoon, evening, or night) of the check-in. The recommended preference (i.e. with whom the check-in was shared) is randomly selected among *Only Me*, *Friends*, *Friends of friends*, and *Public*. For each recommendation, we asked the participants if they would like to use the recommended preference in the future visit to the place at the certain time. To collect more realistic decisions about the recommendations, we have used real check-in data to make recommendations rather than using locations they have never visited.

Final questionnaires

In the third part of the experiment, we collected data on the participants' subjective factors when using our recommenders; specifically we collected data on *quality*, *satisfaction*, and *concern*. Participants were presented with seven questionnaires. The first six evaluated the *quality* and *satisfaction* of the various recommenders. When answering the questionnaires for a particular recommender, participants were shown with their choices to the recommendations from this recommender made in the second part of the experiment, to remind them with their decisions (Figure 3). One final questionnaire measured overall *concern* with the recommenders.

After answering all the questionnaires, in the last step, participants were given opportunities to provide free-text comments of their opinions and suggestions about our recommenders.

RESULTS

99 participants completed our experiment. Table 2 shows their demographic information.

Table 1. Demographic Information

Category	Options	Participants(%)	Facebook(%)
Gender	Female	63	51
	Male	37	49
Age	18-24	74	21
	25-34	24	27
	35-44	2	20
	45-54	0	16
	55+	0	16

Table 2. The demographics of our experiment (Participants) compared with the overall UK Facebook over-18 user population (Facebook). The Facebook data were taken from the Facebook Adverts Manager in October 2015.

Overall, we find a negative effect of *concern* on *acceptance*. For the objective factors of recommendations, the level of openness of the recommended location-privacy preferences and the contexts of the recommendation both influence people's acceptance. The source of crowdsourcing does not influence their acceptance significantly.

Analytical approaches

In our study, we consider *trust*, *quality*, *satisfaction*, and *concern* as the subjective factors. Each of these factors is evaluated using various question items in the questionnaires. Before we evaluate their effects on *acceptance*, we need to first establish their validity. We do this using Confirmatory Factor Analysis (CFA), which can establish both convergent and discriminant validity. Convergent validity ensures that the question items in the same questionnaire measure the same factor, while discriminant validity ensures that two different questionnaires measure two different factors. To maintain convergent validity, question items with low loadings (i.e. the R^2 value) may be removed until the average variance extracted (AVE¹) from each factor is larger than 0.5. To maintain discriminant validity, if two factors are highly correlated, one of them will be removed. By applying CFA we can refine the answers to our questionnaires and increase the validity of the factors in our experiment.


To find out whether there are effects between different factors, we need to propose several hypotheses about these effects and test them to discover significant effects. Structural Equation Modeling (SEM) is an integrative modeling technique that tests all proposed hypotheses simultaneously. By applying SEM, we can combine and analyse the detected effects in an integrative structure that allows us to link all the detected effects together. In our experiment, we use Lavaan [24] to implement both the CFA and SEM analyses.

Privacy concerns lower acceptance of recommendations

We are interested in the effects of *concern*, *trust*, and conditions on *quality*, *satisfaction*, and *acceptance*.

¹The AVE for a given factor is the average of the R^2 values of the factor's question items.

Recommendation 8 (of 30)



We see that you have visited School of Computer Science.

The **Facebook-friends** recommender suggests

You can share your location with

Public
Anyone on or off Facebook

when you are at School of Computer Science on a weekday afternoon.

The next time that you visit School of Computer Science on a weekday afternoon, would you accept the **Facebook-friends** recommender's location-sharing setting?

Yes No

Continue

Figure 2. Each participant in our experiment was presented with 30 recommendations made by our 3 recommenders (10 recommendations from each recommender), and asked if they would accept the recommendation.

Here are all your choices to the recommendations made by the **same-location** recommender

Where	When	To Whom	Your Choice
Odeon Cinema	weekday, evening	friends of friends	no
Wuhan University	weekday, evening	everyone	yes
Parker House	weekday, afternoon	friends of friends	yes
Edinburgh Castle. Edinburgh, Scotland	weekday, evening	everyone	yes
Duke's Corner	weekday, afternoon	friends of friends	no
Overgate	weekday, afternoon	only me	yes
Vic St Andrews	weekend, noon	only me	no
School of Computer Science	weekday, afternoon	only me	yes
Baxter Park	weekday, evening	only me	yes
School of Computer Science	weekend, morning	friends	no

Please answer the following questions to tell us how you think about the **recommendation quality**. (1/6)

1. I like the location-sharing choices that were made by the system.

Strongly disagree Disagree Neutral Agree Strongly agree

2. The recommendations fitted my location-privacy preferences.

Strongly disagree Disagree Neutral Agree Strongly agree

3. The recommended location-sharing choices were well-chosen.

Strongly disagree Disagree Neutral Agree Strongly agree

4. The recommended location-sharing choices were relevant.

Strongly disagree Disagree Neutral Agree Strongly agree

5. The system recommended too many bad location-sharing choices.

Strongly disagree Disagree Neutral Agree Strongly agree

6. I didn't like any of the recommended location-sharing choices.

Strongly disagree Disagree Neutral Agree Strongly agree

7. The recommendations I accepted were "the best among the worst".

Strongly disagree Disagree Neutral Agree Strongly agree

Continue

Figure 3. Questionnaires were used to collect data on perceived recommendation quality (*quality*), satisfaction of choices (*satisfaction*), and concerns about our system (*concern*). Participants were also presented with their choices as a reminder.

Construct	Question items	R^2	AVE
Quality	I like the location-sharing choices that were made by the system.	0.858	0.575
	The recommendations fitted my location-privacy preferences.	0.789	
	The recommended location-sharing choices were well-chosen.	0.822	
	The recommended location-sharing choices were relevant.	0.440	
	The system recommended too many bad location-sharing choices.	0.469	
	I didn't like any of the recommended location-sharing choices.	0.328	
Satisfaction	The recommendations I accepted were "the best among the worst".	0.321	0.520
	I like the recommendations that I've accepted.	0.510	
	Some of my chosen location-sharing choices could become part of my default location-privacy settings.	0.506	
Concern	I would recommend some of the chosen location-sharing choices to others/friends.	0.544	0.602
	I'm afraid that the system discloses private information about me.	0.470	
	The system invades my privacy.	0.861	
	I feel confident that the system respects my privacy.	0.586	
	I'm uncomfortable providing private data to the system.	0.524	
	I think the system respects the confidentiality of my data.	0.571	

Table 3. Results of Confirmatory Factor Analysis (CFA). Question items with low R^2 values are removed in the refined results. The general trust to technology (*trust*) is removed because it only has two question items to keep its AVE greater than 0.5. Both the convergent validity and the discriminant validity of our model hold.

Table 3 shows the refined results of CFA after eliminating low-loading question items. *Trust* only has two question items to make its AVE larger than 0.5 and so it is eliminated from our later SEM analysis, as each factor needs at least 3 question items. Our model's convergent validity holds because the AVEs are greater than 0.5 and its discriminant validity holds too because all the correlations between two different factors are lower than the square roots of the AVEs of both the factors.

After refining the factors, we apply SEM to our refined factors by adding relationships between them. All answers to the question items are modeled as ordinal variables. We use the *same-location* recommender as the baseline condition and introduce two dummy variables *friends* and *similar* to represent the conditions of the *Facebook-friends* recommender and the *similar-people* recommender.

To avoid missing any significant relationships in the SEM model we analyse every possible relationship between factors: for *quality*, we study $quality \sim concern + similar + friends$; for *satisfaction*, we study $satisfaction \sim concern + similar + friends + quality$; and for *acceptance*, we study $acceptance \sim concern + quality + similar + friends + satisfaction$.² The fit³ of our SEM model is adequate: $\chi^2_{125} = 483.67, p < 0.001$; *root mean squared error of approximation (RMSEA)* = 0.098; *Comparative Fit Index (CFI)* = 0.977; *Turker – Lewis Index (TLI)* = 0.972.

Figure 4 shows four significant ($p < 0.001$) effects in our SEM analysis. *Concern* (SC) has a negative effect on both the *quality* (SSA) and the *satisfaction* (EXP), which means

that people with higher privacy concerns about our recommenders perceive lower quality of our recommendations and are less satisfied about their choices. *Quality* acts as a mediator between *concern* and *acceptance* (INT), influencing *acceptance* positively. Another positive effect is from *quality* to *satisfaction*.

Our experimental results indicate that people's **privacy concerns** about location-privacy preference recommendations play an important role when they use such recommenders. We regress all of the "Neutral" answers in our *concern* questionnaire into a baseline *concern* and regress all of the participants' answers to their *concern* factors. We find that 44% of these participants have higher *concern* than the baseline. Their privacy concerns about using our recommenders have a negative influence on their acceptance (mediated by their perceived quality) and their satisfactions about their choices (both directly related and mediated by their perceived quality). Location-privacy preferences are inherently sensitive because they contain both location information and people's privacy settings (i.e. whether they share, and if so, to whom they share the locations). Hence it is not surprising that *concern* influences both *acceptance* and *satisfaction*.

We do not find significant effects from *friends* (OSA) or *similar* (OSA) on *quality*, *satisfaction*, nor *acceptance*. In other words, our participants do not perceive the qualities of the recommended location-privacy preferences to be different even when we claimed that they were made from different crowdsourcing sources.

Openness and context affect acceptance

In addition to the conditions (i.e. the source of crowdsourcing) that were controlled in our experiment, we consider other factors that may influence participants' acceptance of recommendations. These include the level of openness of the recommended location-privacy preferences, and the contexts of the recommendations.

²*Trust* was eliminated at the CFA stage.

³The cut-off values of good fit proposed by Hu and Bentler [13] are: *CFI* > 0.96, *TLI* > 0.95, *RMSEA* < 0.05. However, Kenny et al. [16] recommend not computing the RMSEA for models with small degree of freedom and small sample sizes.

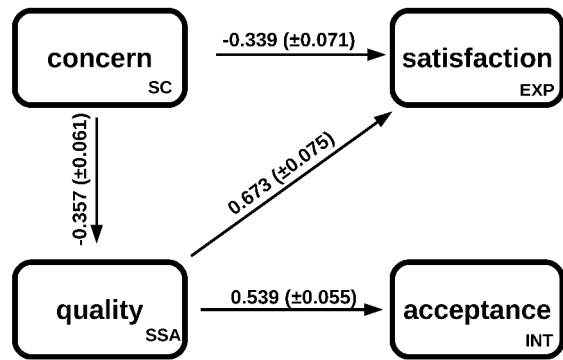


Figure 4. The structured equation modeling (SEM) results. $p < 0.001$ for all coefficients. Numbers above arrows mean the β – weight (\pm standard error) of the effect. Standard deviation = 1. The *concern* has negative effects on *acceptance* (moderated by *quality*) and *satisfaction* (directly and moderated by *quality*).

We evaluate the different acceptance of the recommended preferences for different levels of openness (i.e. with whom the recommender suggests a location is shared). Figure 5 shows the distribution of decisions for levels of openness from very low (sharing with only myself) to high (sharing with the public). The preference with the highest openness has the lowest acceptance, which means people are less likely to accept highly open location-privacy preference recommendations. Nevertheless, the preference with the lowest openness (sharing to only me) also has low acceptance. This means that people not only consider privacy issues, but also care about the benefits (e.g. social needs) from sharing their locations. For example, one of our participants’ feedback includes:

- “... if i (*sic*) would only share something to ‘only me’, then why would i (*sic*) share at all?”

When the benefits are guaranteed (i.e. sharing to someone), the recommendations of low openness location-privacy preferences get higher acceptance.

To study context, we manually analyse the categories given by Facebook of all the locations in our recommendations and merge similar categories with each other. By this means, we have four categories: *Entertainment*, *Residential*, *School/University/Library*, and *Transport*. For the time dimension, we use the five time slots from our previous work [29]: *morning* (07:00–11:59), *noon* (12:00–13:59), *afternoon* (14:00–16:59), *evening* (17:00–20:59), and *night* (21:00–06:59).

Figure 6 shows that for different contexts, our participants have different acceptance of the recommended location-privacy preferences (two-way ANOVA to examine the interaction effect of time and location: $F = 2.039, df = 12, p < 0.05$). In all of the time slots, recommendations for *School/University/Library* locations are most accepted. Since we advertised our experiment through university mail lists and university Facebook groups, we believe that the majority




Figure 5. Distribution of all participants’ acceptance (the percentage of accepted recommendations) of different recommended location-privacy preferences (sharing to whom). The least open preference (*only me*) and the most open preference (*everyone*) are least accepted. For the sharing-preferences (i.e. *friends*, *friends of friends*, and *everyone*), the less open preference is more accepted. (ANOVA: $F = 33.45, df = 3, p < 0.001$)

of our participants are university students. This implies that our participants accept our recommendations mostly in the contexts where they conduct regular daily lives. The *Transport* category experiences the lowest acceptance of recommendations in the three time slots including *morning*, *noon*, and *afternoon*. Comments made by participants may explain this:

- “... for example, I was at the airport. This informs all Facebook users that I will be away for potentially a longer period of time than usual and could put myself at greater risk of property theft etc. ...”
- “... I think a better recommender could consider sharing a place by how regularly you go there or how far from where you normally are it is ie how exotic it is.”

It appears that the **regularity** of contexts and the **potential risks** caused by the false positive location-privacy preference recommendations (i.e. sharing the locations which people do not want to share) in different contexts have influences on people’s acceptance. We plan to study these questions in future research.

DISCUSSION

We find several effects from both objective and subjective factors on people’s acceptance of location-privacy preference recommendations in our user study. We combine these effects together in Figure 7.

As shown in Figure 7, our aim is to design a system where people have both high acceptance of the recommended location-privacy preference and high satisfaction of their choices. These are listed in the left box that represents users’ subjective factors. Since we cannot directly manipulate these subjective factors, a system designer must measure the influential objective factors (the right box) and adjust recommendation strategies to making recommendations that people are more likely to accept. Our experimental results indi-

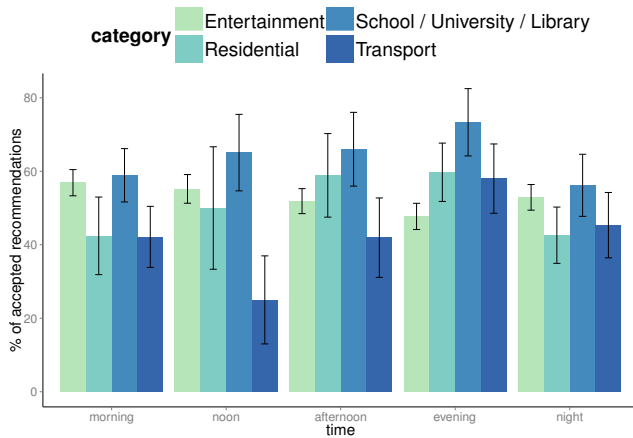


Figure 6. Participants’ acceptance of the recommendations made for different location categories. For each time slot, our participants have the highest acceptance of recommended location-privacy preferences in the School/University/Library category, which is the most regular context for them (two-way ANOVA: $F = 2.039, df = 12, p < 0.05$).




Figure 7. The effects of the subjective factors (left side) of users and the objective factors (right side) of location-privacy preference recommendations. The arrows in solid line mean the detected effects in our experimental results. The dash line means a potential effect from an unknown objective factor (marked as ?) to concern. Recommender system designers can only control the objective factors on the right side to influence the subjective factors on the left side.

cate that the contexts of the recommendations and the level of openness of the recommendations, which can be measured by recommender designers, will influence the *acceptance*. We therefore suggest that these two factors should be taken into account in the implementation of location-privacy preference recommenders. In addition, we found the effect from the *concern* on the left side, but we do not know the factors (marked as ? in Figure 7) on the right side that can influence the *concern*, which is our future work. We now discuss our recommendations in further detail.

Be cautious about the level of openness

The first effect we find is the openness of the recommended location privacy preferences. To our surprise, the lowest-openness recommendations do not receive the highest acceptance. In fact, the acceptance of the lowest-openness recommendation is the second lowest and the highest-openness recommendation has the lowest acceptance. This means that people are not likely to give up the benefits of sharing their

locations for complete location-privacy, nor vice versa. This suggests that location-privacy recommenders should be careful with such extreme recommendations. When extreme recommendations are made, additional information, such as explanations or request for consent, might be necessary to increase people’s acceptance. Meanwhile, overexposing recommendations would decrease people’s acceptance. To avoid this happening, enabling people to control the maximum openness that they allow the recommender to make is also a possible solution. This is supported by our participants’ feedback:

- “If the system had a ‘never share publicly’ option that would work best for my preferences. ...”
- “There should be a ‘maximum exposure’ option ...”

In addition, customised recommendations that can provide finer-grained openness might also be beneficial:

- “..., however I would like some more customization. ...”
- “The recommenders should take into account the preferences I’ve set in the past.”
- “Would need to learn a bit more about my own preferences as well as aggregating those from other sources to be useful for me.”

An outstanding research challenge for customisation is how to make recommendations by using different types of customised preferences because people may have different ways to categorise and name the recipients of their location.

Recommendations should be context-aware

The second effect we find is from the contexts for the recommendations. We observed that our participants have higher acceptance of recommended location-privacy preferences in their most regular context than others. Combined with our post-experiment feedback, we postulate that the **regularity** of and the **potential risks** (e.g. being away from home for a long time) of overexposure in different contexts may also influence people’s acceptance. In future work, we plan to investigate these factors in more detail. But our current results imply that system designers may wish to let users choose in which contexts they want to use the recommender, or tune recommendations to context. We also suggest that recommenders in other domain such as mobile application recommendations may be context-aware [3], since there are evidences from existing research [4] indicate that people’s mobile application usage is highly dynamic with contexts

Recommenders must consider privacy and security

Our experimental results demonstrate that 44% of our participants have privacy concerns when using our recommenders and this concern has a negative effect on their acceptance of the recommended location-privacy preferences and their satisfactions about their choices. People may decide that providing their privacy preferences to a recommender is risky due to its higher sensitivity than other preferences. This type of concern is the *concern* in our experiment. Thus it is necessary to implement location-privacy preference recommenders in a privacy-aware fashion [29].

In addition to privacy, users may also be concerned about security. Recommenders are vulnerable to a series of attacks due to their open structure. One example is the shilling attack [10], which means the malicious users might be able to inject fake profiles and preferences to influence the outcome. Attacks may also be passive, such as the inference attack [5], which means the attackers might be able to infer the location-privacy preferences of certain target users by passively observing the change of the recommendations with the help of some auxiliary information. Additional work is needed to investigate whether these vulnerabilities would contribute to people's concerns and decrease their likelihood of contributing their data to the system. Accordingly, it is important to find solutions to detect and alleviate these attacks.

LIMITATIONS

In our experiment, we find the effects from both the subjective and objective factors on people's acceptance of location-privacy preference recommendations. That said, we note some limitations in our experiment that we hope to overcome in our future work.

First, to ensure the participants were familiar with the contexts of the recommendations, for each participant, we only used their own location check-in history as the contexts for making recommendations. Compared with asking them to consider the recommendations hypothetically in some places they have never been, the answers through our method are more likely to reflect their true decisions. This means, however, that we failed to evaluate their acceptance of the recommended location-privacy settings when they actually enter new places. In future work, we hope to deploy our recommender in real world environments and evaluate people's acceptance of recommended location-privacy preferences *in situ*.

A second limitation is that we used a random recommendation generator rather than real recommenders. This was done to keep the objective quality of the recommendations the same for all participants, since the performance of real recommendations for a specific user could be influenced by the quality of their own preference. By using the same random recommender for all participants, we could ensure that all the changes we observed were due to the changes that we controlled. As a consequence, the overall performance of the recommendations was impacted inevitably. Since the source of crowdsourcing was shown to not have an influence on people's acceptance when the recommenders have the same performance, it would be useful to investigate which crowdsourcing source has the best objective performance when deploying a system in the real world.

CONCLUSIONS

In this paper, we conduct an online user experiment to investigate which objective and subjective factors influence people's acceptance of location-privacy preference recommendations. Our results show that the openness of the recommended location-privacy preferences and the contexts of the recommendations both have effects. Meanwhile, people's privacy

concerns about using our recommenders have a negative influence on their acceptance of the recommended location-privacy preferences and the satisfaction about their choices. Compared with existing research focusing on the objective performance (i.e. recommendation accuracy), our work sheds light on the parameters that system designers might want to control and the issues which need to be solved when designing location-privacy preference recommenders in order to make them more acceptable to users.

Our future plans are to investigate solutions, both technical and social, for alleviating the security and privacy issues that affect users' concerns about privacy preference recommenders. We believe that this will result in systems that are both technically better, and better accepted by users.

REFERENCES

1. D. Anthony, T. Henderson, and D. Kotz. Privacy in Location-Aware Computing Environments. *IEEE Pervasive Computing*, 6(4):64–72, Oct. 2007. doi:10.1109/MPRV.2007.83.
2. G. Bigwood, F. Ben Abdesslem, and T. Henderson. Predicting Location-Sharing Privacy Preferences in Social Network Applications. In *Proceedings of the First Workshop on recent advances in behavior prediction and proactive pervasive computing (AwareCast)*, 2012.
3. M. Böhmer, G. Bauer, and A. Krüger. Exploring the design space of context-aware recommender systems that suggest mobile applications. In *2nd Workshop on Context-Aware Recommender Systems (CARS)*, 2010.
4. M. Böhmer, B. Hecht, J. Schöning, A. Krüger, and G. Bauer. Falling asleep with Angry Birds, Facebook and Kindle: A Large Scale Study on Mobile Application Usage. In *Proceedings of MobileHCI 2011*, pages 47–56, 2011. doi:10.1145/2037373.2037383.
5. J. a. Calandrino, A. Kilzer, A. Narayanan, E. W. Felten, and V. Shmatikov. “You might also like:” Privacy risks of collaborative filtering. *Proceedings of the IEEE Symposium on Security and Privacy*, pages 231–246, 2011. doi:10.1109/SP.2011.40.
6. S. Das, T. H.-J. Kim, L. A. Dabbish, and J. I. Hong. The Effect of Social Influence on Security Sensitivity. In *Proceedings of SOUPS'14*, July 2014.
7. C. Dong, H. Jin, and B. P. Knijnenburg. Predicting Privacy Behavior on Online Social Networks. In *Proceedings of ICWSM*, pages 91–100, 2015.
8. L. Fang and K. LeFevre. Privacy wizards for social networking sites. In *Proceedings of the WWW conference*, page 351, 2010. doi:10.1145/1772690.1772727.
9. S. Furnell. Managing privacy settings: lots of options, but beyond control? *Computer Fraud & Security*, 2015(4):8–13, Apr. 2015. doi:10.1016/S1361-3723(15)30027-0.

10. I. Gunes, C. Kaleli, A. Bilge, and H. Polat. Shilling attacks against recommender systems: a comprehensive survey. *Artificial Intelligence Review*, 42(4):767–799, Dec. 2014. doi:10.1007/s10462-012-9364-9.
11. J. L. Harman, J. O’Donovan, T. Abdelzaher, and C. Gonzalez. Dynamics of human trust in recommender systems. In *Proceedings of RecSys 2014*, pages 305–308, 2014. doi:10.1145/2645710.2645761.
12. B. Henne, C. Kater, and M. Smith. On Usable Location Privacy for Android with Crowd-Recommendations. In *Proceedings of TRUST 2014*, pages 74–82, 2014. doi:10.1007/978-3-319-08593-7_5.
13. L. Hu and P. M. Bentler. Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*, 6(1):1–55, Jan. 1999. doi:10.1080/10705519909540118.
14. L. Hutton and T. Henderson. An architecture for ethical and privacy-sensitive social network experiments. *ACM SIGMETRICS Performance Evaluation Review*, 40(4):90–95, Apr. 2013. doi:10.1145/2479942.2479954.
15. Q. Ismail, T. Ahmed, A. Kapadia, and M. K. Reiter. Crowdsourced Exploration of Security Configurations. In *Proceedings of CHI 2015*, pages 467–476, 2015. doi:10.1145/2702123.2702370.
16. D. A. Kenny, B. Kaniskan, and D. B. McCoach. The Performance of RMSEA in Models With Small Degrees of Freedom. *Sociological Methods & Research*, 44(3):486–507, Aug. 2015. doi:10.1177/0049124114543236.
17. B. P. Knijnenburg and H. Jin. The persuasive effect of privacy recommendations for location sharing services. In *SIGHCI 2013 Proceedings*, 2013. doi:10.2139/ssrn.2399725.
18. B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, and C. Newell. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*, 22(4-5):441–504, Oct. 2012. doi:10.1007/s11257-011-9118-4.
19. J. Lindqvist, J. Cranshaw, J. Wiese, J. Hong, and J. Zimmerman. I’m the mayor of my house: Examining Why People Use foursquare - a Social-Driven Location Sharing Application. In *Proceedings of CHI 2011*, pages 2409–2418, 2011. doi:10.1145/1978942.1979295.
20. M. Madejski, M. Johnson, and S. M. Bellovin. A study of privacy settings errors in an online social network. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*, pages 340–345, Mar. 2012. doi:10.1109/PerComW.2012.6197507.
21. S. M. McNee, J. Riedl, and J. A. Konstan. Making recommendations better: an analytic model for human-recommender interaction. In *Proceedings of the CHI 2006 extended abstracts*, pages 1103–1108, 2006. doi:10.1145/1125451.1125660.
22. K. D. Naini, I. S. Altingovde, R. Kawase, E. Herder, and C. Niederée. Analyzing and Predicting Privacy Settings in the Social Web. In *Proceedings of UMAP 2015*, pages 104–117, 2015. doi:10.1007/978-3-319-20267-9_9.
23. P. Pu, L. Chen, and R. Hu. A user-centric evaluation framework for recommender systems. In *Proceedings of RecSys 2011*, pages 157–164, 2011. doi:10.1145/2043932.2043962.
24. Y. Rosseel. lavaan: An R Package for Structural Equation Modeling. *Journal of Statistical Software*, 48(2), 2012. doi:10.18637/jss.v048.i02.
25. N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao. Understanding and capturing people’s privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, Aug. 2009. doi:10.1007/s00779-008-0214-3.
26. J. Staiano, N. Oliver, B. Lepri, R. de Oliveira, M. Caraviello, and N. Sebe. MoneyWalks: A Human-Centric Study on the Economics of Personal Mobile Data. In *Proceedings of UbiComp 2014*, pages 583–594, 2014. doi:10.1145/2632048.2632074.
27. E. Toch. Crowdsourcing privacy preferences in context-aware applications. *Personal and Ubiquitous Computing*, 18(1):129–141, Dec. 2012. doi:10.1007/s00779-012-0632-0.
28. J. Xie, B. P. Knijnenburg, and H. Jin. Location sharing privacy preference: analysis and personalized recommendation. In *Proceedings of IUI 2014*, pages 189–198, 2014. doi:10.1145/2557500.2557504.
29. Y. Zhao, J. Ye, and T. Henderson. Privacy-aware Location Privacy Preference Recommendations. In *Proceedings of Mobiquitous 2014*, pages 120–129, 2014. doi:10.4108/icst.mobiquitous.2014.258017.
30. A. H. Zins and U. Bauernfeind. Explaining Online Purchase Planning Experiences with Recommender Websites. In *Information and Communication Technologies in Tourism 2005*, pages 137–148, 2005. doi:10.1007/3-211-27283-6_13.