DRAFT only. To appear in the proceedings of the UMAP 2011 Workshop on Decision Making and Recommendation Acceptance Issues in Recommender Systems, Girona, Spain.

Exploring the Effects of Feed-forward and Feedback on Information Disclosure and User Experience in a Context-Aware Recommender System

Bart Knijnenburg[1], Alfred Kobsa[1], Simon Moritz[2], Martin A. Svensson[2],

[1] Department of Informatics, University of California, Irvine, USA
{Bart.K, Kobsa}@uci.edu
[2] Ericsson Research, Ericsson AB, Stockholm, Sweden
{simon.moritz, martin.a.svensson}@ericsson.com

Abstract. When disclosing information to a recommender system, users need to trade off its usefulness for receiving better recommendations with the privacy risks incurred through its disclosure. Our paper describes a series of studies that will investigate the use of feed-forward and feedback messages to inform users about the potential usefulness of their disclosure. We hypothesize that this approach will influence the user experience in several interesting ways.

Keywords: Recommender systems, privacy, information disclosure, context-aware recommenders, accuracy, user experience, satisfaction.

1 Introduction

Recommender systems for mobile applications need to provide immediate benefit to users, or else they may discontinue using them [1][2]. Many recommender systems, however, give adequate recommendations after an extensive period of use only [3]. Context-aware recommender systems use context data to overcome this new-user problem. Previous recommender systems have used location, system usage behavior, demographics, and implicit feedback [4].

Some users may feel uneasy providing such potentially privacy-sensitive information to the system [5][6]. From a privacy perspective, it is better to let users decide themselves whether or not they want to disclose some piece of information [7]. Research shows that a large majority of people is willing to trade off privacy for personal benefits [8]. However, users often have a hard time making an informed decision because they lack knowledge about their benefit from providing the information to the system and its consequences for their privacy [9][10].

Recent studies on users' election of privacy settings in an IM client [11] and a Facebook application [12] informed participants about the privacy decisions made by their friends (or in [12]: all other users). This "feed-forward" message facilitated "social cues" [13]; participants were slightly more likely to conform to the social norm in setting their privacy preferences. The current paper implements this idea in the field of recommender systems, and presents several extensions.

2 Other Types of Feed-forward and Feedback

While previous work has considered the impact of social cues only, we plan to investigate a variety of feed-forward messages that can help users make educated information disclosure decisions (Table 1). Wang and Benbasat [14] showed that providing feed-forward about the usefulness of the piece of information to be disclosed increased users' trust in the recommender system. Berendt and Teltzrow [15] suggest that providing such information might also increase the amount of disclosure. We propose a similar feed-forward message, which promises users that the recommendations will improve by a certain amount if they disclose a certain piece of information. The social cues and usefulness promises can be combined in a feed-forward message that tells users what percentage of other users received better recommendations after disclosing the information in question. Like previous work, in our studies the numbers in the feed-forward messages, which reflect the level of influence of the messages, will not be based on real data but will rather be random within given ranges (Table 2).

Table 1. Different types of feed-forward messages to be investigated in our studies.

Type of feed-forward	Message to user
None	(no message)
Social	"XX% of our users gave us/allowed us to use…"
Usefulness	"The recommendations will be about XX% better when you give us/ allow us to use…"
Social usefulness (combined)	"XX% of our users received better recommendations when they gave us/allowed us…

Table 2. Different levels of influence that will be used in the feed-forward messages.

Level	Percentage
Low	A random number between 5% and 25%
Moderate	A random number between 40% and 60%
High	A random number between 75% and 95%

Whereas participants in previous studies set their privacy settings once, we propose a system in which users can decide to change the amount and type of information they disclose. Users may base this decision on two pieces of feedback: the quality of the recommendations they receive, and a reflection of the information they are disclosing ('detailed profile inspection'). The effect of the quality of the recommendations on the amount of disclosure is unclear. In 'conversational' recommenders, where users incrementally disclose information, users tend to disclose more information if they see that this increases the recommendation quality [16]. This effect may however not occur in a system where most of the disclosure is at the beginning of the interaction.

For those types of disclosure that accumulate information over time, the user may initially not be aware of the exact extent of the disclosure. It is therefore assumed to be good privacy practice to allow users to inspect the 'profile' that the system has gathered over time [17]. Such detailed profile inspection may assist the user in deciding whether to change her information disclosure settings (see Table 3).

Table 3. Different levels of profile inspection (feedback).

Type of feedback	Implementation
Shallow	Shows the types of information being disclosed (e.g. "app usage"), but no specific information (e.g. the usage frequency)
Detailed	Shows the types of information being disclosed, as well as a detailed record of this information

3 Information Elicitation and the User Experience

In our proposed system, the amount of disclosure has a direct impact on the quality of the recommendations, and consequently on users' satisfaction with the system. Information disclosure is thus a tradeoff between usefulness of disclosure and protection of privacy. Providing users with information can nudge users into over-protecting or under-protecting their privacy. If users are lured into over-protection, their satisfaction may decrease because the recommender may not have enough information to generate accurate recommendations. If users are lured into under-protection, they may later feel that their privacy was compromised.

Merely looking at users' level of disclosure paints a one-sided picture; the complex nature of users' interaction with the system warrants an integrative, user-centric approach. Based on Knijnenburg et al. [18], we hypothesize that several factors mediate the effect of feed-forward, feedback and disclosure on user experience: perceived privacy threat, perceived amount of control over the system, trust, and perceived quality of the recommendations. Table 4 and Figure 1 show the hypothesized effects.

Table 3. Different levels of profile inspection (feedback).

Topic	Hypotheses
Feed-forward and feedback	The different types of feed-forward messages (H1) and levels of usefulness (H2) have a different impact on the initial amount of disclosure. The profile inspector (H3) and recommendation quality (H4) influence the change in disclosure. The profile inspector increases the perceived control over the privacy settings (H5), which increases the trust in the system (H6), which in turn causes a (negative) change in the level of disclosure (H7).
Privacy concerns and privacy threats	Users' privacy concerns decrease the amount of initial disclosure (H8) and cause a (negative) change in disclosure (H9). The amount of initial disclosure (H10), change in disclosure (H11), and users' privacy concerns (H12) influence the perceived privacy threat.
Recommendation quality and choice satisfaction	The amount of disclosure (H13) and change in disclosure (H14) influence the perceived recommendation quality, which in turn influences the satisfaction with the installed apps (H15)
System satisfaction and system use	The perceived privacy threat (H16), perceived recommendation quality (H17), and perceived control over the settings (H18) influence the system satisfaction, which is in turn related to the extent of system use (H19)

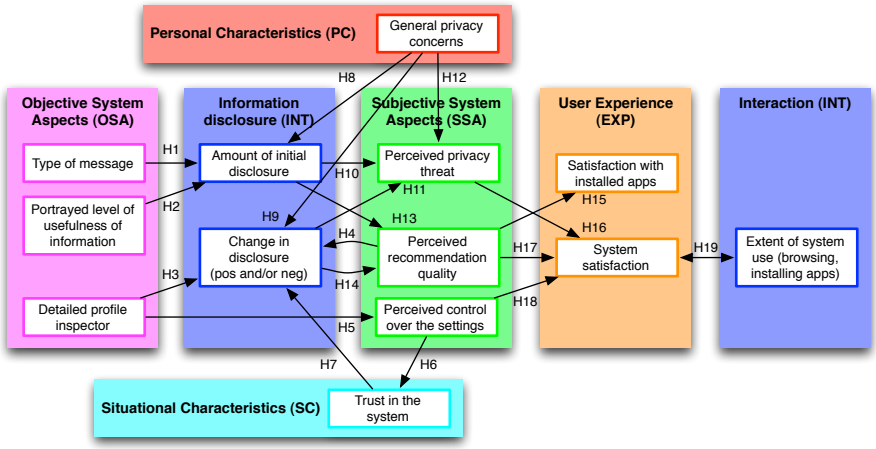

Fig. 1. A visual representation of the hypothesized effects in our studies.

4 Proposed Studies

We propose a series of studies that implement and test our feed-forward and feedback mechanisms in an app recommender system developed by Ericsson [3] with the working title "Applause" (Figure 2). The system asks users to disclose their location (Figure 2, screen 1), current app usage (e.g. app download, forwarding to friends, usage frequency, location, and time of day; screen 2), app browsing behavior in the system (screen 3), and demographics (e.g. age, income, occupation; screen 4). To guarantee that our findings are both comprehensive and statistically valid, we propose a variety of studies: qualitative user interviews, an online questionnaire, a highly controlled experiment with a system mockup, and a field test with real users of the real system.

4.1 Qualitative Study

The goal of the qualitative study is to get an in-depth insight into how users trade off the benefits of disclosing information with the threats that this poses to their privacy. 20-30 participants will be recruited, and given the opportunity to use the current version of Applause (without feed-forward and feedback) for at least a week.

Participants are asked to elaborate on their experience with the system. They are also asked about their phone usage, technological expertise, and privacy concerns. After that, they are shown different mockups of information disclosure screens (Figure 2, screen 1-4). Screens will display different types of feed-forward messages, as well as different levels of influence. For each screen, users are asked if they would disclose the information or not, and to elaborate on their decision. Participants are also shown the different levels of profile inspection (screen 7-8), and asked for their comments. The goal of this study is to explore users' reactions to changes on each dimension.

Interview responses will be analyzed using grounded theory analysis [19], which models relationships between concepts (i.e. type of message and privacy concerns). Models of each participant are compared to identify similarities and conflicts. The interviews are conducted in three batches, so that insights from the first analysis can influence the questions asked in the second batch of interviews. Finally, an integrated model is constructed, and interesting deviations from this model are highlighted.

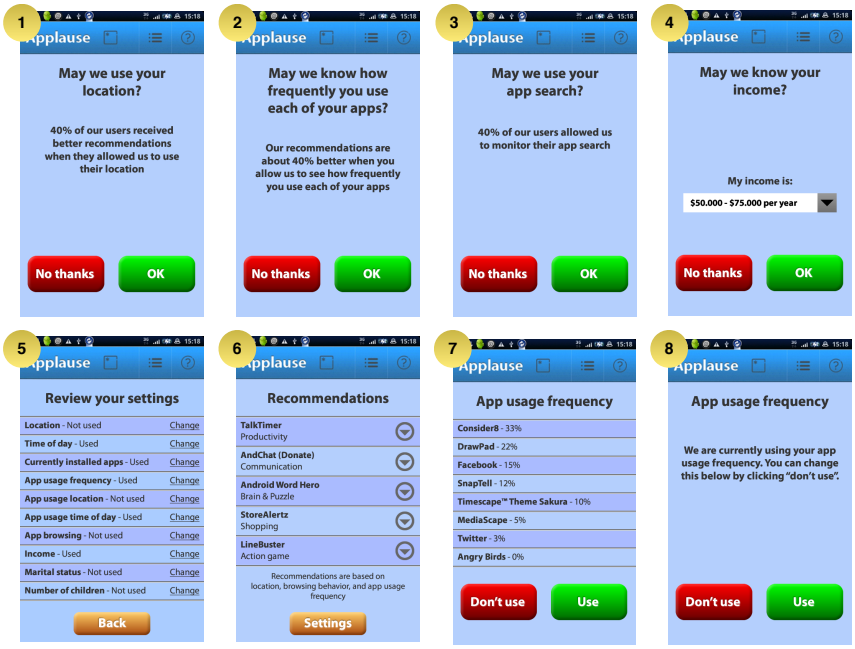

Fig. 2. Mockups of feed-forward and feedback conditions in proposed app recommender.

4.2 Online Survey

The online survey has the same goal as the quantitative study, but its results will be based on a larger sample (150-200 participants) and will have a quantitative character, allowing statistical validation of the results. This study also pre-tests the questionnaires that will be used in subsequent studies. Whereas participants in the quantitative study were asked to compare different types of feed-forward and feedback, the quantitative study presents each user with only one type of feed-forward message and one type of profile inspection. This results in 2x4 between-subjects conditions. The type of requested information and the level of influence are manipulated within subjects.

Participants are first asked several questions about their phone usage, technological expertise and privacy concerns. They are then randomly assigned to an experimental condition and shown mockups of information disclosure screens, with the feed-forward message and level of influence corresponding to this condition (Figure 2, screen 1-4). For each screen, participants are asked to disclose this information or not. They are also shown one of the two profile inspectors (screen 7-8), and asked if they

want to change their disclosure. Finally, they are asked several questions about the perceived privacy threat that this system poses, their perceived control over their profile, their satisfaction with a system that would use these features, and their intention to use that system and to recommend it to a friend.

Structural equation modeling will be used to extract relevant subjective concepts from the questionnaire responses and determine relationship between the experimental conditions and these concepts. The hypothesized effects that can be tested are a subset of the ones displayed in Figure 1; specifically, due to the setup of the experiment we cannot test H1-H3, H5-H12, H16, H18, and H19. As participants in this study are not interacting with the system, use can only be measured as an intention.

4.3 Fake Recommendation Experiment

Whereas the first two studies ask participants about their intended use of the system, the two experiments described in this and the next section consider actual system use. Research has shown that privacy attitudes and behaviors do not always align [20][21]. The fake recommendation experiment uses a semi-functional mockup of the recommender system that does not provide real recommendations (i.e., every participant receives the same recommendations), thereby controlling for the effects that would normally be mediated by the recommendation quality. Because the system is used only once, the different types of profile inspection will not be considered in this study. Type of feed-forward is again manipulated between subjects, and type of information and level of influence within subjects. The design of the study resembles [11]. The main difference to this work is that we test different types of messages.

100-150 participants are first asked several questions about their phone usage, technological expertise and privacy concerns. They then interact with the system in one cycle. The system first asks them to disclose information (Figure 2, screen 1-4), where each screen shows a feed-forward message that corresponds to the randomly selected condition and the randomly selected level of influence. Then the system provides the (fake) recommendations (Figure 2, screen 6). Finally, participants are asked about the perceived privacy threat posed by this system, the perceived quality of the recommendations, their perceived control over their own profile, their satisfaction with the system, and their intention to use the system if it would be available.

Structural equation modeling will be used to statistically test the relationships between the experimental conditions, the disclosure behavior, the subjective system aspects, and the user experience. The following hypotheses in Figure 1 will be tested: H1, H2, H8, H10, H11, H13, H15, H16, H17 and H19. Note that participants use the system only once, so “extent of system use” can only be measured as an intention.

4.4 Field Experiment

The field experiment uses the fully operational app recommender. The study will sample 350 to 500 participants from existing users of the Applause system. Participants are shadowed over a period of time (in which they will be allowed to change their disclosure), and receive real recommendations based on their disclosure. The

type of feed-forward message and the type of profile inspection are manipulated between subjects (leading to 2x4 conditions), and the type of requested information and the level of influence are manipulated within subjects.

Participants are first asked several questions about their phone usage, technological expertise and general privacy concerns. Consequently, they interact with the system repeatedly for a period of two weeks. Their initial interaction will be the same as in the fake recommendation experiment. However, after the first information elicitation screens (Figure 2, screen 1-4), participants are asked to review their settings (screen 5) before moving on to the recommendations (screen 6). Participants are encouraged to revisit the recommendation screen throughout the study period. They will also be informed that they can return to the review screen to change their disclosure. When changing their disclosure, some participants are aided by a detailed profile inspector (screen 7), while others will only see a global profile inspector (screen 8).

The system logs participants' information disclosure and system usage (browsing recommendations, installing recommended apps). After two weeks, participants are asked several questions about the perceived privacy threat that this system poses, the perceived quality of the recommendations, and their satisfaction with the system and the apps they installed that were recommended by the system. Structural equation modeling will be used to evaluate all hypotheses in Figure 1.

5 Conclusion and Future Work

Employing the user-centric framework for recommender system evaluation in [18], this paper applies (and extends) recent findings on information disclosure [11] to the field of recommender systems. Information disclosure is important for the proper operation of most recommender systems, and privacy issues are specifically salient in context-aware recommenders, where disclosure moves beyond the traditional elicitation of preferences. All proposed studies include "pretend" elements. Even the final study uses a "fake" feed-forward message (e.g., the expected usefulness of a certain piece of information is not actually calculated). More research needs to be done to find 'real' metrics of information usefulness (e.g. the expected amount of change, or increase in accuracy, in the recommendations when providing the information).

References

1. Xiao, B., Benbasat, I.: E-commerce Product Recommendation Agents: Use, Characteristics, and Impact. MIS Quarterly. 31, 137--209 (2007).
2. Kobsa, A.: Privacy-Enhanced Web Personalization. In: Brusilovsky, P., Kobsa, A., and Nejdl, W. (eds.) The Adaptive Web: Methods and Strategies of Web Personalization. pp. 628--670. Springer, Heidelberg (2007). DOI: 10.1007/978-3-540-72079-9_21
3. Davidsson, C., Moritz, S.: Utilizing Implicit Feedback and Context to Recommend Mobile Applications from First Use. IUI 2011 Workshop on Context-awareness in Retrieval and Recommendation. pp. 19--22. ACM, New York (2011). DOI: 10.1145/1961634.1961639
4. Goy, A., Ardissono, L., Petrone, G.: Personalization in E-Commerce Applications. In: Brusilovsky, P., Kobsa, A., and Nejdl, W. (eds.) The Adaptive Web: Methods and Strategies

of Web Personalization. p. 485--520. Springer, Heidelberg (2007). DOI: 10.1007/978-3-540-72079-9_16
5. Chellappa, R.K., Sin, R.G.: Personalization Versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. Information Technology and Management. 6, 181--202 (2005). DOI: 10.1007/s10799-005-5879-y
6. Teltzrow, M., Kobsa, A.: Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study. In: Karat, C.-M., Blom, J., and Karat, J. (eds.) Designing Personalized User Experiences for eCommerce. pp. 315--332. Kluwer, Dordrecht (2004). DOI: 10.1007/1-4020-2148-8_17
7. Solove, D.: The Digital Person: Technology and Privacy in the Information Age. New York University Press, New York (2004).
8. Acquisti, A.: Privacy in Electronic Commerce and the Economics of Immediate Gratification. 5th ACM Conference on Electronic Commerce. pp. 21--29. ACM Press, New York (2004). DOI: 10.1145/988772.988777
9. Brodie, C., Karat, C.-M., Karat, J.: Creating an E-Commerce Environment Where Consumers are Willing to Share Personal Information. In: Karat, C.-M., Blom, J.O., and Karat, J. (eds.) Designing Personalized User Experience for eCommerce. pp. 185--206. Kluwer, Dordrecht (2004). DOI: 10.1007/1-4020-2148-8_11
10. Kobsa, A., Teltzrow, M.: Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data Sharing Behavior. In: Martin, D. and Serjantov, A. (eds.) PET 2005. LNCS, vol 3424, pp. 329--343. Springer, Heidelberg (2005). DOI: 10.1007/11423409_21
11. Patil, S., Page, X., Kobsa, A.: With a Little Help from My Friends. In: ACM 2011 Conference on Computer Supported Cooperative Work. pp. 391--394. ACM Press, New York (2011). DOI: 10.1145/1958824.1958885
12. Besmer, A., Watson, J., Lipford, H.R.: The Impact of Social Navigation on Privacy Policy Configuration. In: 6th Symposium on Usable Privacy and Security. ACM Press, New York (2010). DOI: 10.1145/1837110.1837120
13. Dourish, P., Chalmers, M.: Running out of space: Models of Information Navigation. Short paper presented at HCI (1994).
14. Wang, W., Benbasat, I.: Recommendation Agents for Electronic Commerce: Effects of Explanation Facilities on Trusting Beliefs. Journal of Management Information Systems. 23, 217--246 (2007). DOI: 10.2753/MIS0742-1222230410
15. Berendt, B., Teltzrow, M.: Addressing Users' Privacy Concerns for Improving Personalization Quality. In: Mobasher, B., and Sarabjot, A.S. (eds.) LNCS, vol. 3169, pp. 69--88 (2005). DOI: 10.1007/11577935_4
16. Knijnenburg, B.P., Willemsen, M.C., Hirtbach, S.: Receiving Recommendations and Providing Feedback: The User-Experience of a Recommender System. In: Buccafurri, F. and Semeraro, G. (eds.) EC-Web 2010. LNBIP, vol. 61, pp. 207--216. Springer, Heidelberg (2010). DOI: 10.1007/978-3-642-15208-5_1
17. Kay, J.: Stereotypes, Student Models and Scrutability. In: Gauthier, G., Frasson, C., and VanLehn, K. (eds.) Intelligent Tutoring Systems. pp. 19--30. Springer, Heidelberg (2000).
18. Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H., Newell, C.: Explaining the user experience of recommender systems, http://db.tt/JG7079A.
19. Glaser, B.: The discovery of grounded theory: Strategies for qualitative research. Aldine Transaction, New Brunswick (2006).
20. Spiekermann, S., Grossklags, J.: E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior. 3rd ACM Conference on Electronic Commerce. pp. 38--47. ACM Press, New York (2001). DOI: 10.1145/501158.501163
21. Tsai, J.Y., Egelman, S., Cranor, L.F., Acquisti, A.: The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. Information Systems Research. 21, (2010). DOI: 10.1287/isre.1090.0260