

Inspectability and Control in Social Recommenders

Bart P. Knijnenburg
Department of Informatics
University of California, Irvine

bart.k@uci.edu

John O'Donovan
Department of Computer Science
University of California, Santa Barbara

jod@cs.ucsb.edu

Svetlin Bostandjiev
Department of Computer Science
University of California, Santa Barbara

alex@cs.ucsb.edu

Alfred Kobsa
Department of Informatics
University of California, Irvine

kobsa@uci.edu

ABSTRACT

Users of social recommender systems may want to inspect and control how their social relationships influence the recommendations they receive, especially since recommendations of social recommenders are based on friends rather than anonymous “nearest neighbors”. We performed an online user experiment (N=267) with a Facebook music recommender system that gives users control over the recommendations, and explains how they came about. The results show that inspectability and control indeed increase users’ perceived understanding of and control over the system, their rating of the recommendation quality, and their satisfaction with the system.

Categories and Subject Descriptors

H.1.2. [Models and principles]: User/Machine Systems—*software psychology*; H.4.2. [Information Systems Applications]: Types of Systems—*decision support*; H.5.2 [Information Interfaces and Presentation]: User Interfaces—*evaluation/methodology, interaction styles, user centered design*

General Terms

Measurement, Design, Experimentation, Human Factors, Theory.

Keywords

Social recommender systems, human-computer interaction, usability, user experience, user interfaces, control, inspectability, visualization, explanations, novelty, understandability, satisfaction

1. INTRODUCTION

Collaborative recommender systems compare users’ preferences to those of all other users, and recommend items that are liked by those users who have similar preferences [21]. Social recommenders limit the set of other users to your friends, thereby leveraging personal connections [37, 39, 52]. We suspect that users of social recommenders may not be satisfied with only a static list of recommendations. Rather, they may want to *inspect* and *control* the way in which the system uses their social network to select this list of recommendations, for at least two reasons:

1. Users seem to appreciate it when recommender systems explain their recommendations [9, 12, 21, 47, 48, 50]. In social recommenders, where users know the people on which the recommendations are based, the system can provide such explanation by showing how the overlap between one’s preferences and those of one’s friends resulted in a set of recommendations. Such a “recommendation graph” increases the *inspectability* (or transparency [9, 47]) of a system, which could have a positive effect on users’ experience [44].
2. Users seem to appreciate *control* in their interaction with recommender systems [30, 35]. Recommenders have to somehow gather users’ preferences, and different types of preference elicitation methods provide different levels of control [6, 29]. In a recommender system that leverages social networks such as Facebook, the system can use users’ “likes” to construct a preference model, and the overlap with their friends’ “likes” to compute recommendations. However, users may want some control over this process, because they may not like each item equally well, or they may value a friend’s preferences beyond (or short of) the amount of mutual overlap in “likes”. Users may therefore want to give additional (or lower) weight to some of the items and/or some of their friends.

Although this reasoning may seem intuitive, little research has been done to establish the effect of inspectability and control on the users’ experience with social recommender systems ([17, 18] are notable exceptions). This paper describes the results of an online user experiment (N=267) with a Facebook music recommender in which we independently manipulated the level of inspectability and control. The results show that the versions of our system that offer high inspectability and control indeed provide a better user experience. The structural model in which we present these results allows us to explain *why* inspectability and control are important qualities of social recommender systems.

2. RELATED WORK

Before discussing the experiment we first survey related work on the effects of inspectability and control in recommender systems in general and in social recommenders specifically. We also dis-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

RecSys '12, September 9–13, 2012, Dublin, Ireland.

Copyright 2012 ACM 978-1-4503-1270-7/12/09...\$15.00.

This work was partially supported by the U.S. Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053 and by NSF grant IIS-1058132; The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ARL, NSF, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

cuss related work on the user-centric evaluation of recommender systems, and on personal characteristics of the user that may influence our results.

2.1 Inspectability

Many aspects of the explanation process have been studied in the recommender systems literature. For example, an earlier approach in [22] discusses a simple glass-box paradigm that provides only minimal information to the user. In this paper our notion of inspectability is similar to Tintarev and Masthoff's concept of transparency in [44], which is to "explain how the system works", and it is treated separately from control in our study. [46] also introduces the concept of scrutability with the (more interactive) aim to "allow users to tell the system it is wrong", while Czarkowski and Kay [10] examine scrutability and control as separate mechanisms, in the context of a student model application. Kay and Lum [26] also focus on scrutability, but in terms of providing explanations of why individual elements and relations in the underlying model have particular values.

Herlocker [19] argues that explanation provides transparency, "exposing the reasoning behind a recommendation". The reasoning and insight into the recommendation process exposed by an explanation interface can also increase the *inspectability* of the system as a whole. Tintarev and Masthoff [48] show that explanations make it easier to judge the quality of recommendations. Consequently, such explanations increase users' trust in the recommendations and, in turn, the perceived competence of the system ([9, 12], see also [19, 50]). Sinha and Swearingen [44] demonstrate that users rate systems that provide detailed information about items as more useful and easier to use.

In the realm of social recommenders, Groh et al. [18] present a study that outlines the "extensive need" for explanation, and Gretarsson et al. [17] present a small-scale study of an explanation interface, finding that the explanation process has a positive effect on satisfaction with recommendations.

2.2 Control

Researchers have implemented various aspects of control in recommender systems, ranging from simple preference elicitation at recommendation time [7] to more complex iterative processes such as dynamic critiquing which allows users to tweak ordered numerical attributes during the recommendation phase [6, 32, 33, 42]. More recent work [17, 37, 39] discusses interactive graphical representations of the recommendation process, to enable control over both item- and user-level preferences in collaborative recommender systems.

Multiple studies highlight the benefits of interactive interfaces that support control over the recommendation process. In a general comparison of user-controlled versus static recommendation interfaces, McNee et al. [35] found that study participants preferred user-controlled interfaces because these systems "best understood their tastes". McNee et al. also showed that participants had higher retention rates with the controlled interface. Knijnenburg et al. [30] found that controllable recommendations are typically deemed more varied than automatic ones, and Willemsen et al. [51] show that diversifying recommendations can be useful to overcome choice overload.

Other common methods of control include rating items [14, 41] and assigning weights to item attributes [20, 31]. Research shows that the choice between these different methods has a substantial impact on the user experience [6, 29, 31].

2.3 Social Recommender Systems

Previous studies have attempted to explain recommendations by showing the link between the recommendations and the "nearest neighbors" on which they are based [21, 35, 47]. An interesting aspect of social recommenders is that recommendations are based on users' similarity with their *friends* rather than a set of anonymous nearest neighbors. In effect, social recommenders can leverage users' acquaintance with the source of the recommendation, which instantly attaches a wealth of established social information to the recommendations that can be further explored and exploited in the processes of inspection and control [18].

Specifically, we hypothesize that visualizing the link between recommendations and the nearest neighbors on which they are based increases the inspectability of social recommenders beyond regular recommenders, because the neighbors are known.

Furthermore, in social recommenders one could allow the user to not only rate items, but also their friends [5]. This would give certain friends additional (or less) weight beyond the weight computed based on preference overlap. Arguably, this method allows users to indicate how much they "trust" their friends' preferences in the recommendation domain. Several researchers have investigated this idea of assigning trust scores to friends in collaborative recommenders, through explicit mechanisms such as in Golbeck's FilmTrust system [15] which can support propagation of trust scores around a network of peers, and through automated mechanisms for modeling trust such as [11, 38]. Several recent studies have extended these ideas to prediction of personality, and by derivation, behavior of a user within the system [1] in terms of both trust and distrust [16].

In this paper we take a step beyond existing work on social recommenders by explicitly testing the effect of leveraging users' knowledge about their friends to improve inspectability and control. By allowing a user to inspect and control the elements of trust on which the recommendations are based, we can gain an understanding of the effect of inspectability and control on the user experience with the recommender system.

2.4 User Experience

In order to determine the impact of inspectability and control, we need to measure users' *experience* with the recommender system. Specifically, we are not only interested in the quality of the recommendations, but also in users' satisfaction with the system as a whole. Moreover, we need to consider concepts like understandability and perceived control to explain how inspectability and control influence the users' experience.

Knijnenburg et al. [30] developed a framework for user-centric evaluation of recommender systems through user experiments. It describes how certain manipulations (in our case: inspectability and control) influence subjective system aspects (i.e. understandability, perceived control and recommendation quality), which in turn influence user experience (i.e. system satisfaction).

2.5 Personal Characteristics

The framework also allows us to include the effect of *personal characteristics* on users' experience. In social recommenders, the degree of trustfulness (or "trusting propensity") could play an important role, because in order to accept the recommendations, users will need to trust their friends' preferences [15, 38]. Users who are not trustful usually want to take matters in their own hands and therefore demand more control over the system [49].



Instructions

- By clicking on the boxes below, you can see how your likes are linked to your friends, and how your friends are linked to the recommendations.
- Please carefully inspect the visualization and the recommendations by clicking on the boxes below.
- When you are done, click "Next".



Figure 1. The TasteWeights system as used in the online user experiment. This is the inspection phase of the “full graph” condition. Users can click on items, friends and recommendations to see the links between them. The inspection phase of the “list only” condition shows the rightmost list (recommendations) only.

Another personality trait that can have an influence on users’ experience is choice persistence, a characteristic that divides users into satisficers and maximizers. Satisficers will stop their search when they encounter an item that meets their criteria, while maximizers continue their efforts until they find the best possible option [43]. Maximizers may thus show more appreciation for systems that allow extensive control and inspection (although we do not find evidence for this in [29]).

Finally, users’ domain knowledge can also have a significant impact on their experience with a recommender system. Kamis and Davern show that domain experts perceive personalized recommenders as less useful than novices [24], while other researchers have consistently found that experts have a higher appreciation for the recommendations as well as the recommender system itself [3, 30, 51]. In order to be satisfied with a recommender, however, domain experts want more control over their recommendations than novices [13, 29, 40, 45]. Novices, on the other hand, require a simple and understandable recommender system [31], and they may even prefer to give up control in return for simplicity [29].

3. ONLINE USER EXPERIMENT

The related work shows that inspectability and control have a positive influence on users’ experience with recommender systems in general, and we suspect that these benefits may be even more pronounced in social recommenders. A user study is needed to investigate the nature and extent of these benefits, and the factors that influence them. We therefore conducted an online user experiment with 267 participants employing a modified version of the TasteWeights [4, 17] social recommender.

3.1 System

The TasteWeights system recommends new artists/bands based on the music “likes” of the user and her Facebook friends.

3.1.1 Recommendation algorithm

The TasteWeights system calculates its recommendations in two steps. First, weights are computed for each friend based on their similarity to the user. Specifically, the similarity of the user to a friend is given by the overlap in music “likes” between them, as defined by Pearson’s correlation coefficient:

$$W_{friend_i} = \frac{TWCI_{user, friend_i}}{\sqrt{TWI_{user}^2 \cdot TWI_{friend_i}^2}}$$

where $TWCI_{x,y}$ is the total weight of the items (“likes”) that users x and y have in common, and TWI_x is the total weight of items liked by user x . As Facebook users can only “like” artists/bands without specifying how much, item weights are initialized to 0.5.

Once all friend weights are computed, recommendations are generated by assigning weights to all friends’ music items (excluding the items that the user already “likes”):

$$W_{rec_i} = \sum_{friend_j \text{ likes } rec_i} W_{friend_j}$$

Here, the weight of a recommendation i is the sum of the weights of all friends that like i . The recommendations are displayed in decreasing order of recommendation weight.

3.1.2 Inspectability and Control

In terms of inspectability, the TasteWeights system displays a graph (Figure 1) that shows the users’ items, their friends, and the recommendations. By clicking at the graph, the connections between these entities can be explored. The system also shows a short description for each recommended band/artist with a link to their LastFM information page.

The system allows two types of control over the recommendations: users can adjust the weights of their items (initialized at 0.5) and their friends (initially weighted by similarity). Changing the weight of an item will influence the similarity scores, and thus the

recommendation weights. Changing the weight of a friend will add or subtract a proportion from that friend’s similarity score, and thus also influence the recommendation weights.

3.2 Study setup

3.2.1 Experimental conditions

The original TasteWeights system allows users to interactively inspect and control the recommendation graph (i.e. change the weights and inspect the graph simultaneously and iteratively). However, to investigate the effects of inspectability and control independently, we let participants in our experiment interact with the system in two stages: a control stage and an inspection stage.

In the control stage, participants are assigned to one of three conditions (Figure 2): they either skip the control stage altogether (the “no control” condition), they are asked to adjust the weights of the items they “like” (the “item control” condition), or they are asked to adjust the weights of their friends (the “friend control” condition). Our primary interest is to see how these control conditions compare against the no control condition, but we are also interested in differences between the two control conditions.

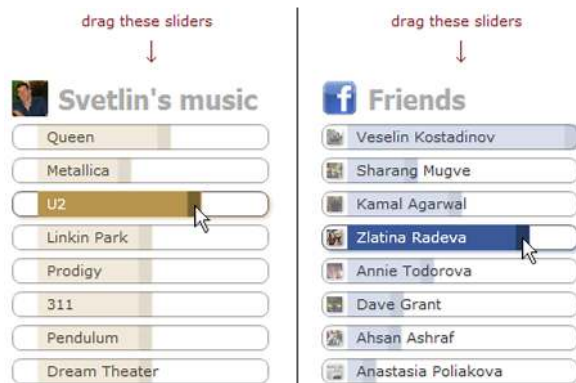


Figure 2. The control phase of item control (left) and friend control (right) conditions.

In the inspection stage, participants are assigned to one of two conditions: the system either shows only the list of recommendations (the “list only” condition), or the full recommendation graph (the “full graph” condition). To give each participant a comparable experience, we limited the number of music likes and friends to be considered by the recommender to 10 each (with a minimum of 5 each). The number of recommendations was fixed to 10.

3.2.2 Participants and Procedure

267 participants were recruited using Craigslist and Amazon Mechanical Turk¹. Only adult Facebook users living in the U.S. were allowed to participate. 156 participants were female, and 130 were between the ages of 18 and 25, 114 between 25 and 40, and 23 older than 40². In order to provide a meaningful experience, we only allowed users to participate if their recommendation graph would show at least 5 music “likes”, showing overlap with at least 5 friends, and resulting in at least 10 recommendations. Denied

¹ In [27] we found no substantial differences between these two participant populations.

² These numbers reflect the general Facebook population, with a slight underrepresentation of the older demographic. See <http://bit.ly/insidefacebook-20100104>.

participants were given the suggestion to populate their Facebook profile with more music “likes” and then try again.

Eligible participants were then asked to answer 15 questions about their personal characteristics (music expertise, trusting propensity and choice persistence). Questions were statements to which participants could agree or disagree on a five point scale. They subsequently completed the control stage (unless they were assigned to the “no control” condition), in which they were asked to adjust the weights of either their items or their friends (depending on the control condition). Next, they completed the inspection stage, where they were asked to carefully inspect the list of recommendations or the recommendation graph (depending on the inspectability condition). Finally, they were asked to indicate whether they already knew each of the recommended band/artist or not, and subsequently to rate the recommendation on a 5-star scale. Users were encouraged to click on the provided LastFM link to improve their judgment of unknown bands/artists. After the experiment, participants answered another 29 questions about their user experience. Full questionnaires can be found in [28].

3.2.3 Questionnaires

The answers to the 44 questionnaire items were submitted to a confirmatory factor analysis³ with categorical indicators and a weighted least squares estimator, estimating 7 factors:

- *Music expertise*: 4 items adapted from [3], e.g. “Compared to my peers I listen to a lot of music.”, α : .74, AVE: .627
- *Trusting propensity*: 3 items adapted from [29], e.g. “In general, people really do care about the well-being of others.”, α : .80, AVE: .657
- *Understandability*: 3 new items, e.g. “The recommendation process is clear to me.”, α : .92, AVE: .877
- *Perceived control*: 4 items adapted from [29], e.g. “Compared to how I normally get recommendations, TasteWeights was very limited.” (reversed), α : .84, AVE: .643
- *Perceived recommendation quality*: 5 items adapted from [30], e.g. “I liked the bands/artists recommended by the TasteWeights system.”, α : .90, AVE: .738
- *System satisfaction*: 7 items adapted from [30], e.g. “I can find better music using TasteWeights.”, α : .92, AVE: .708
- *Familiarity with recommenders*: 2 new items, e.g. “I am familiar with online recommender systems.”, α : .86, AVE: .794

10 questionnaire items were excluded from the analysis due to low communality, high cross-loadings and/or high residual correlations. Additionally, 5 items on choice persistence (taken from [36]) failed to converge to a stable factor. For the remaining factors the values of Cronbach’s α and average variance extracted (AVE) were high⁴, indicating convergent validity. Moreover, the square root of the AVE is higher than the factor correlation for all factors, indicating discriminant validity.

3.3 Results

We subjected the 7 factors, the experimental conditions, and selected behaviors to structural equation modeling, which simultaneously fits the factor measurement model and the structural rela-

³ Factor analysis and structural equation modeling as applied in this paper are explained in Appendix A of [30].

⁴ For alpha, >.70 is acceptable, >.80 is good, >.90 is excellent. AVE should be >.50 for convergent validity.

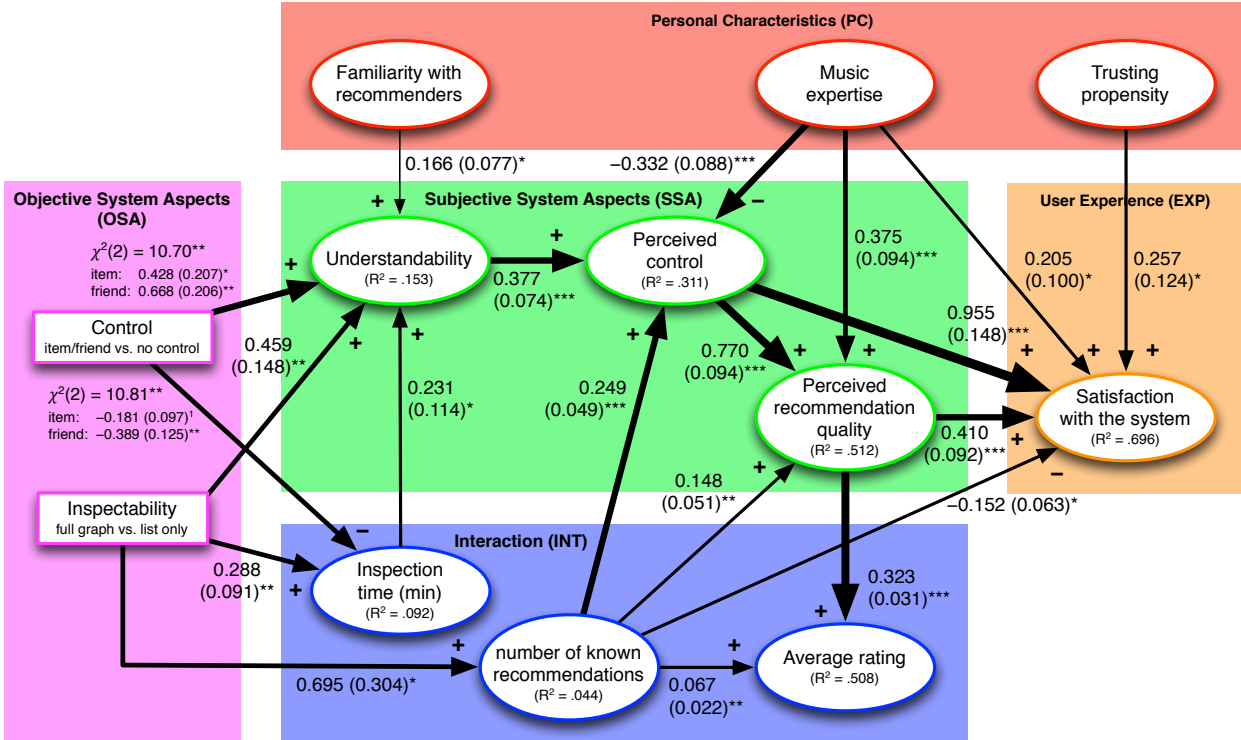


Figure 3. The structural equation model for the data of the experiment. Significance levels: * $p < .001$, ** $p < .01$, 'ns' $p > .05$. R^2 is the proportion of variance explained by the model. Numbers on the arrows (and their thickness) represent the β coefficients (and standard error) of the effect. Factors are scaled to have an SD of 1.**

tions between factors and other variables. The model (Figure 3) has a good⁵ model fit: $\chi^2(537) = 639.22, p < .01; RMSEA = 0.027, 90\% CI: [0.017, 0.034], CFI = 0.993, TLI = 0.992$.

3.3.1 Subjective Experience

The model shows that the inspectability and control manipulations each have an independent positive effect on the understandability of the system: the full graph condition is more understandable than the list only condition, and the item control and friend control conditions are more understandable than the no control condition (see also Figure 4a). Understandability is in turn related to users' perception of control, which is in turn related to the perceived quality of the recommendations. The perceived control and the perceived recommendation quality finally determine participants' satisfaction with the system (for the marginal effects of control and inspectability on these factors, see Figure 4b,c,d).

3.3.2 User Behavior

There exist additional effects of inspectability and control on understandability, which are mediated by the inspection time (the amount of time users take to inspect the recommendations, see Figure 4e). In the full graph condition, participants take more time to inspect the recommendations (about 7.3 seconds more), and this results in an additional increase of understandability. For the two control conditions, however, the inspection time is shorter (about 10.9 seconds less in the item control condition and about

23.3 seconds less in the friend control condition), which counters the positive effect on understandability.

In the full graph condition, participants indicate that they already know more of the recommendations than in the list only condition (see Figure 4f). In turn, the more recommendations the participant already knows, the higher is the perceived control and perceived recommendation quality, but the lower is the satisfaction.

The perceived recommendation quality and the number of known recommendations determine the average rating participants give to the recommendations. The marginal effects of the inspectability and control manipulations on the average rating (Figure 4g) indicate that the ratings in the item control condition are somewhat lower (mean: 3.146) than the no control condition (mean: 3.267), whereas the ratings in the friend control condition are somewhat higher (mean: 3.384). The difference between the two control conditions is small but significant ($p = .031$).

3.3.3 Personal Characteristics

Participants who are familiar with recommenders find the system more understandable. Participants with music expertise perceive less control over the system, but perceive a higher recommendation quality and system satisfaction. Finally, trusting propensity influences participants' satisfaction with the system.

4. Discussion

Based on the results of our experiment, we can describe in detail how the benefits of inspectability and control in social recommenders come about. We can also describe these results in the light of users' personal characteristics. Finally, we can provide some preliminary suggestions on the relative effectiveness of controlling items versus friends.

⁵ A model should not have a non-significant χ^2 , but this statistic is regarded as too sensitive [2]. Hu and Bentler [23] propose cut-off values for other fit indices to be: $CFI > .96, TLI > .95$, and $RMSEA < .05$, with the upper bound of its 90% CI below 0.10.

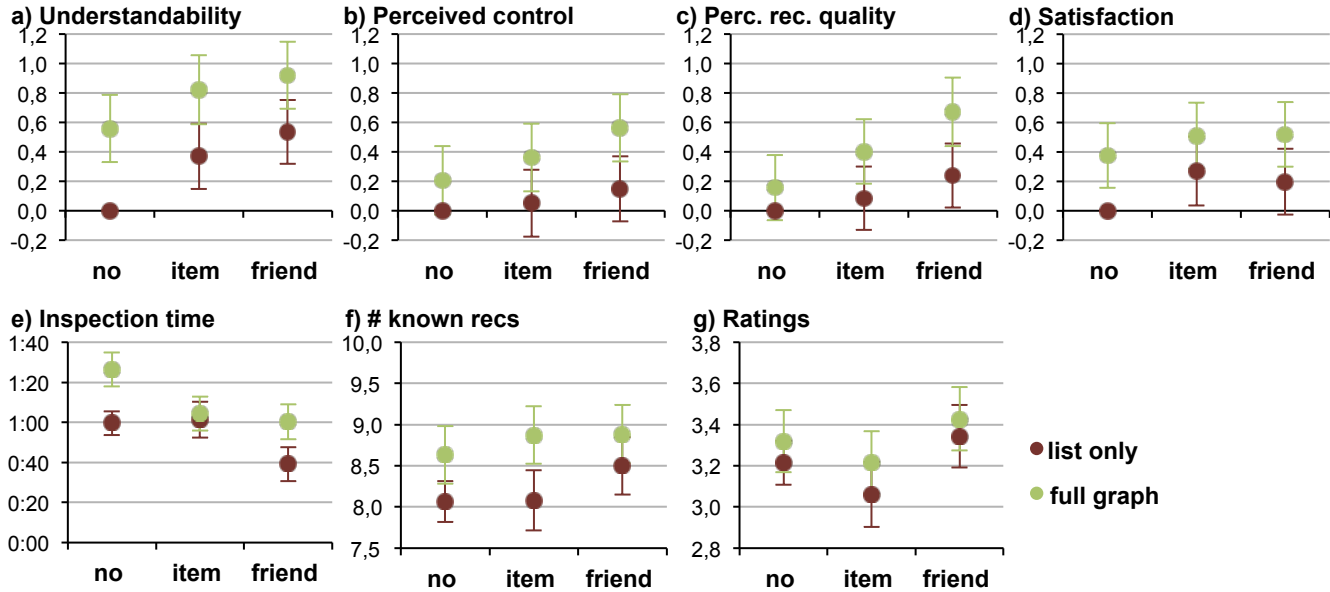


Figure 4. Marginal effects of inspectability and control on the subjective factors (top) and on behaviors (bottom). For the subjective factors, the effects of the “no control, list only” condition is set to zero, and the y-axis is scaled by the sample standard deviation.

4.1 Inspectability and Control

Both inspectability and control have a positive effect on the user experience, primarily because an inspectable and controllable recommender system is easier to understand. The increased understandability causes users to feel more in control over the system, and this in turn increases the perceived quality of the recommendations, also indicated by increased ratings. Finally, the higher perceived control and recommendation quality cause users to be more satisfied with the system.

Inspectability works partially due to a direct effect on understandability, and partially due to its influence on user behavior. Specifically, users take more time for inspection in the “full graph” condition (which increases understandability), and users in this condition already know more of the recommendations (which increases perceived control and recommendation quality, but decreases system satisfaction). The effect of inspectability on the number of recommendations that the participant already knows may seem counterintuitive, because the inspectability conditions do not influence the actual recommendations. However, in the “full graph” condition users can see which friends are connected to the recommendations, and this may allow users to *recognize* more of the recommendations as already known (e.g. “I remember John playing this band’s album for me”)⁶.

Arguably, this recognition effect is an important aspect of inspectability, because knowing recommendations may raise users’ trust in the recommender [8, 44]. In our experiment, known recommendations increase users’ perceived control (total effect: $\beta = 0.372, p = .001$) and the perceived recommendation quality (total effect: $\beta = 0.389, p = .002$). On the other hand, known recommendations are less useful, as they contain no novelty, which explains the decrease in system satisfaction (McNee et al. [34] show that users are happy with a set of recommendations as long as it con-

tains at least one novel item). Despite this negative effect of known items, the total effect of inspectability on system satisfaction is however still statistically significant: $\beta = 0.409, p = .001$.

Item control and friend control result in a more understandable system despite the shorter inspection time (total effects: $\beta = 0.386, p = .063$ and $\beta = 0.578, p = .004$, respectively). Note that although inspection time is shorter, participants in these conditions spend additional time controlling the recommendations.

4.2 Personal Characteristics

Several personal characteristics have an effect on users’ experience when using our system. Trusting propensity has a positive effect on system satisfaction, which may be due to the fact that users with a higher general trusting propensity seem more likely to trust their friends’ music preferences. Arguably, then, trustfulness is an important precondition for a social recommender to work for a user.

Moreover, users with some expertise about music feel less in control, but they view the recommendations and the system itself more positively. Music experts may feel that bands/artists are too crude of a building block for recommendations (for them, bands may have both amazing and terrible albums), which could have caused the reduced perception of control (this effect is consistent with findings in [24]). On the other hand, music experts are more capable of judging the quality of the recommendations, which may be the reason for the increased perceived recommendation quality and satisfaction with the system (these effects are consistent with findings in [3, 30, 51]).

4.3 Which Type of Control?

Besides comparing the control conditions against the “no control” condition, we are also interested in comparing the control conditions against each other, to determine which type of control users prefer. Figure 4 shows that the understandability, perceived control and perceived recommendation quality are consistently higher for the “friend control” condition than for the “item control” condition, but the difference between these two conditions is not sta-

⁶ Conformity bias could be an alternative explanation: “If all my friends know this band, I ought to know it too!”

tistically significant. The only significant difference between the two control conditions is in users' ratings of the recommendations: "friend control" results in higher ratings, but the difference is again very small (3.146 vs. 3.384 on a 5-star scale).

On the other hand, the "friend control" condition results in slightly more known recommendations, which may be one reason why the system satisfaction is also slightly worse in the "friend control" condition. With a note of caution due to the lack of significant results, we can interpret these trends to conclude that the "friend control" condition may give the user more accurate control, but the "item control" condition may lead to a perception of more novel recommendations.

5. CONCLUSION AND FUTURE WORK

Our results show that social recommender systems (and arguably recommender systems in general) indeed benefit from facilities that improve their inspectability and control.

Recommenders that display a recommendation graph result in a better user experience, partially because they encourage users to spend more time on inspecting the recommendations. Moreover, by inspecting the links between friends and recommendations, users get hints about their previous encounters with the recommended items, which allows them to more accurately evaluate the quality of the recommendations. In effect, the recommendation graph increases the understandability, perceived control, perceived recommendation quality, and system satisfaction.

Recommenders that give users control over the item weights and friend weights are more understandable, which leads to higher perceived control, perceived recommendation quality and overall system satisfaction. As to which control method is preferred, the results suggest that giving users control over the friend weights results in slightly higher quality recommendations, but that controlling the item weights may heighten users' perception of recommendation novelty. Arguably, making both control mechanisms available may preserve the best aspects of both methods.

In fact, as we find that the effects of inspectability and control are additive, the best user experience may arise when users can control items and friends simultaneously and directly in the recommendation graph. Such interactive control could arguably result in *scrutability*: allowing users to find and correct mistakes in the recommendation process [25, 47]. In this paper we purposefully disentangled inspectability and control to isolate their respective effects; in future work we purpose to investigate the benefits of scrutability in a fully interactive social recommender.

6. ACKNOWLEDGEMENTS

We thank Nikhil Rao for his extensive help conducting the study and Tobias Höllerer for his feedback throughout the process.

7. REFERENCES

- [1] Adali, S., Escrivá, R., Goldberg, M.K., Hayvanovych, M., Magdon-Ismael, M., Szymanski, B.K., Wallace, W.A. and Williams, G. Measuring behavioral trust in social networks. *2010 IEEE Intl. Conference on Intelligence and Security Informatics* (Vancouver, Canada, 2010), 150–152.
- [2] Bentler, P.M. and Bonett, D.G. Significance Tests and Goodness of Fit in the Analysis of Covariance Structures. *Psychological Bulletin*. 88, 3 (1980), 588–606.
- [3] Bollen, D., Knijnenburg, B.P., Willemsen, M.C. and Graus, M. Understanding choice overload in recommender systems. *Proceedings of the fourth ACM conference on Recommender systems* (Barcelona, Spain, 2010), 63–70.
- [4] Bostandjiev, S., O'Donovan, J. and Höllerer, T. Taste-Weights: A Visual Interactive Hybrid Recommender System. *Proceedings of the sixth ACM conference on Recommender Systems* (Dublin, Ireland, 2012).
- [5] Bourke, S., McCarthy, K. and Smyth, B. Power to the people: exploring neighbourhood formations in social recommender system. *Proceedings of the fifth ACM conference on Recommender systems* (New York, NY, 2011), 337–340.
- [6] Chen, L. and Pu, P. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction*. 22, 1-2 (Oct. 2011), 125–150.
- [7] Chen, L. and Pu, P. *Survey of Preference Elicitation Methods*. Technical Report #IC/2004/67 (2004).
- [8] Cooke, A.D.J., Suján, H., Suján, M. and Weitz, B.A. Marketing the Unfamiliar: The Role of Context and Item-Specific Information in Electronic Agent Recommendations. *Journal of Marketing Research*. 39, 4 (Nov. 2002), 488–497.
- [9] Cramer, H., Evers, V., Ramlal, S., Someren, M., Rutledge, L., Stash, N., Aroyo, L. and Wielinga, B. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction*. 18, 5 (Aug. 2008), 455–496.
- [10] Czarkowski, M. and Kay, J. Bringing Scrutability to Adaptive Hypertext Teaching. *Intelligent Tutoring Systems*. Springer (2000). 423–432.
- [11] DuBois, T., Golbeck, J. and Srinivasan, A. Predicting Trust and Distrust in Social Networks. *IEEE 3rd intl. conference on social computing* (Boston, MA, Oct. 2011), 418–424.
- [12] Felfernig, A. Knowledge-Based Recommender Technologies for Marketing and Sales. *Intl. J. of Pattern Recognition and Artificial Intelligence*. 21, 2 (2007), 333–354.
- [13] Fitzsimons, G.J. and Lehmann, D.R. Reactance to Recommendations: When Unsolicited Advice Yields Contrary Responses. *Marketing Science*. 23, 1 (Jan. 2004), 82–94.
- [14] Gena, C., Brogi, R., Cena, F. and Vernerio, F. The Impact of Rating Scales on User's Rating Behavior. *User Modeling, Adaption and Personalization*. J.A. Konstan, R. Conejo, J.L. Marzo, and N. Oliver, eds. Springer (2011). 123–134.
- [15] Golbeck, J. Generating Predictive Movie Recommendations from Trust in Social Networks. *Trust Management*. K. Stølen, W.H. Winsborough, F. Martinelli, and F. Massacci, eds. Springer (2006). 93–104.
- [16] Golbeck, J., Robles, C. and Turner, K. Predicting personality with social media. *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems* (Vancouver, Canada, 2011), 253–262.
- [17] Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C. and Höllerer, T. SmallWorlds: Visualizing Social Recommendations. *Computer Graphics Forum*. 29, 3 (2010), 833–842.
- [18] Groh, G., Birnkammerer, S. and Köllhofer, V. Social Recommender Systems. *Recommender Systems for the Social Web*. Springer (2012). 3–42.
- [19] Guy, I., Ronen, I. and Wilcox, E. Do you know?: recommending people to invite into your social network. *Proceedings of the 14th international conference on Intelligent user interfaces* (New York, NY, 2009), 77–86.
- [20] Haubl, G. and Trifts, V. Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids. *Marketing Science*. 19, 1 (Jan. 2000), 4–21.
- [21] Herlocker, J.L., Konstan, J.A. and Riedl, J. Explaining collaborative filtering recommendations. *Proc. of the 2000 ACM*

- conference on Computer supported cooperative work (Philadelphia, PA, 2000), 241–250.
- [22] Höök, K., Karlgren, J., Wärn, A., Dahlbäck, N., Jansson, C., Karlgren, K. and Lemaire, B. A glass box approach to adaptive hypermedia. *User Modeling and User-Adapted Interaction*. 6, 2-3 (Jul. 1996), 157–184.
- [23] Hu, L. and Bentler, P.M. Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*. 6, 1 (Jan. 1999), 1–55.
- [24] Kamis, A. and Davern, M.J. Personalizing to product category knowledge: exploring the mediating effect of shopping tools on decision confidence. *37th an. Hawaii intl. conf. on System Sciences* (Big Island, HI, 2004), 10 pp.
- [25] Kay, J. Stereotypes, Student Models and Scrutability. *Intelligent Tutoring Systems*. G. Gauthier, C. Frasson, and K. VanLehn, eds. Springer (2000). 19–30.
- [26] Kay, J. and Lum, A. Ontology-based user modelling for the Semantic Web. *Workshop on Personalisation on the Semantic Web* (Edinburgh, UK, 2005), 15–23.
- [27] Knijnenburg, B.P. and Kobsa, A. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *submitted to the ACM Transactions on Intelligent Interactive Systems, Special Issue on Human Decision Making and Recommender Systems* (2012). http://www.isr.uci.edu/tech_reports/UCI-ISR-12-1.pdf
- [28] Knijnenburg, B.P., Rao, N. and Kobsa, A. *Experimental Materials Used in the Study on Inspectability and Control in Social Recommender Systems*. Technical Report #UCI-ISR-12-4 (2012).
- [29] Knijnenburg, B.P., Reijmer, N.J.M. and Willemsen, M.C. Each to his own: how different users call for different interaction methods in recommender systems. *Proceedings of the fifth ACM conference on Recommender systems* (Chicago, IL, 2011), 141–148.
- [30] Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H. and Newell, C. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*. 22, 4-5 (2012), 441–504.
- [31] Kramer, T. The Effect of Measurement Task Transparency on Preference Construction and Evaluations of Personalized Recommendations. *Journal of Marketing Research*. 44, 2 (May. 2007), 224–233.
- [32] McCarthy, K., Salem, Y. and Smyth, B. Experience-Based Critiquing: Reusing Critiquing Experiences to Improve Conversational Recommendation. *Case-Based Reasoning. Research and Development*. I. Bichindaritz and S. Montani, eds. Springer (2010). 480–494.
- [33] McGinty, L. and Reilly, J. On the Evolution of Critiquing Recommenders. *Recommender Systems Handbook*. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer (2011). 419–453.
- [34] McNee, S.M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S.K., Rashid, A.M., Konstan, J.A. and Riedl, J. On the recommending of citations for research papers. *Proceedings of the 2002 ACM conference on Computer supported cooperative work* (New Orleans, LA, 2002), 116–125.
- [35] McNee, S.M., Lam, S.K., Konstan, J.A. and Riedl, J. Interfaces for Eliciting New User Preferences in Recommender Systems. *User Modeling 2003*. P. Brusilovsky, A. Corbett, and F. Rosis, eds. Springer Heidelberg. 178–187.
- [36] Nenkov, G.Y., Morrin, M., Ward, A., Schwartz, B. and Hullah, J. A short form of the Maximization Scale: Factor structure, reliability and validity studies. *Journal of Judgment and Decision Making*. 3, (June 2008), 371–388.
- [37] O’Donovan, J., Gretarsson, B., Bostandjiev, S., Hollerer, T. and Smyth, B. A Visual Interface for Social Information Filtering. *Intl. Conf. on Computational Science and Engineering* (Vancouver, Canada, 2009), 74–81.
- [38] O’Donovan, J. and Smyth, B. Trust in recommender systems. *Proceedings of the 10th international conference on Intelligent user interfaces* (San Diego, CA, 2005), 167–174.
- [39] O’Donovan, J., Smyth, B., Gretarsson, B., Bostandjiev, S. and Höllerer, T. PeerChooser: visual interactive recommendation. *Proc. of the 26th SIGCHI conf. on Human factors in computing systems* (Florence, Italy, 2008), 1085–1088.
- [40] Pereira, R.E. Optimizing human-computer interaction for the electronic commerce environment. *Journal of Electronic Commerce Research*. 1, 1 (2000), 23–44.
- [41] Pommeranz, A., Broekens, J., Wiggers, P., Brinkman, W.-P. and Jonker, C.M. Designing interfaces for explicit preference elicitation: a user-centered investigation of preference representation and elicitation process. *User Modeling and User-Adapted Interaction*. 22, 4-5 (Mar. 2012), 357–397.
- [42] Pu, P., Chen, L. and Kumar, P. Evaluating product search and recommender systems for E-commerce environments. *Electronic Commerce Research*. 8, 1-2 (2008), 1–27.
- [43] Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K. and Lehman, D.R. Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*. 83, 5 (2002), 1178–1197.
- [44] Sinha, R.R. and Swearingen, K. Comparing Recommendations Made by Online Systems and Friends. *DELOS Workshop on Personalisation and Recommender Systems in Digital Libraries* (2001), 64–67.
- [45] Spiekermann, S. and Paraschiv, C. Motivating human-agent interaction: Transferring insights from behavioral marketing to interface design. *Electronic Commerce Research*. 2, 3 (2002), 255–285.
- [46] Tintarev, N. and Masthoff, J. A Survey of Explanations in Recommender Systems. *Data Engineering Workshop* (Istanbul, Turkey, Apr. 2007), 801–810.
- [47] Tintarev, N. and Masthoff, J. Designing and Evaluating Explanations for Recommender Systems. *Recommender Systems Handbook*. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer (2011). 479–510.
- [48] Tintarev, N. and Masthoff, J. Evaluating the effectiveness of explanations for recommender systems. *User Modeling and User-Adapted Interaction*. 22, 4-5 (Feb. 2012), 399–439.
- [49] Vries, P. *Trust in systems: effects of direct and indirect information*. Thesis Technische Universiteit Eindhoven (2004).
- [50] Wang, W. and Benbasat, I. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems*. 23, 4 (2007), 217–246.
- [51] Willemsen, M.C., Graus, M.P., Knijnenburg, B.P. and Bolten, D. Not just more of the same: Preventing choice overload in recommender systems by offering small diversified sets. *Submitted for journal publication*. (2012).
- [52] Zhao, S., Zhou, M.X., Yuan, Q., Zhang, X., Zheng, W. and Fu, R. Who is talking about what: social map-based recommendation for content-centric social websites. *Proceedings of the fourth ACM conference on Recommender system* (Barcelona, Spain, 2010), 143–150.

