TUDelft

Delft University of Technology

Someone really wanted that song but it was not me!
Evaluating Which Information to Disclose in Explanations for Group Recommendations

Najafian, Shabnam; Inel, Oana; Tintarev, Nava

Someone really wanted that song but it was not me!

Evaluating Which Information to Disclose in Explanations for Group Recommendations

Shabnam Najafian
Delft University of Technology
Delft, the Netherlands
s.najafian@tudelft.nl

Oana Inel
Delft University of Technology
Delft, the Netherlands
O.Inel@tudelft.nl

Nava Tintarev
Delft University of Technology
Delft, the Netherlands
n.tintarev@tudelft.nl

ABSTRACT

Explanations can be used to supply transparency in recommender systems (RSs). However, when presenting a shared explanation to a group, we need to balance users' need for privacy with their need for transparency. This is particularly challenging when group members have highly diverging tastes and individuals are confronted with items they do not like, for the benefit of the group. This paper investigates which information people would like to disclose in explanations for group recommendations in the *music* domain.

CCS CONCEPTS

• **Human-centered computing** → **User studies**; **Empirical studies in HCI**; *Laboratory experiments*; • **Information systems** → **Recommender systems**.

KEYWORDS

Natural Language Explanations, Social Choice-based Aggregation Strategies, Group Recommendations, Privacy

1 INTRODUCTION

The main focus of current RSs is to propose items to *individual* users. However, in many domains (e.g., music) people often consume items in *groups*, rather than individually. One of the reasons recommending to groups is challenging is that different members of the group may have highly diverging tastes. In this context, presenting an explanation of how the system came up with the recommended item(s), can fulfill the explanatory goal of *transparency* and may make it easier for users to accept items they might not like for the benefit of the group [9].

However, when explaining recommendations for groups of users especially with different tastes rather than individuals, additional goals such as *privacy* become relevant as well. These two goals pose

a trade-off between a) understanding why a specific item has been recommended, and b) their need to feel safe, by preserving their private preferences and information from others in the group [3]. This raises the question of *which information an explanation should disclose when displayed to the whole group in a "low consensus"*.

Although there exists many studies on group recommendations, only a few of them focus on *generating explanations* in the context of group recommendations [4, 5, 10]. Besides, to the best of our knowledge, privacy has never been investigated for formulating the explanations of group recommendations. The main contribution of this work is to gain better insights into which information users want to disclose in explanations in the context of group. This is challenging in the group context, as users' need for privacy is likely to conflict with their need for transparency [3].

We (dynamically) generate natural natural language explanations for group music recommendations which can be adapted for three different scenarios and privacy settings. The scenarios allowed us to compare users' privacy preferences for two "low consensus" cases; where either a) the active user (Unhappy user), or b) their acquaintances (Unhappy acquaintances) did not get their preferred song; with a "high consensus" (Baseline), where both the active user and acquaintances get their preferred song.

2 GENERATING NATURAL LANGUAGE EXPLANATIONS

Below, we describe how we generated natural language explanations. One of our requirements is that the user should be able to decide whether to show/ hide different pieces of their personal information in the explanation. Our templates are designed in a way that can flexibly support the addition or removal of three kinds of information: name, rating, personality. For instance, if users decide to, for example, hide the names but show the personality, no names will appear and the corresponding sentence will be anonymized as follow: *"... This decision does not support the preferences of all the group members. However, it supports the preferences of some group members who really want to listen to this song and won't be talked out of it easily"*. These explanations are always generated for a group of three people, with one active user, and their two acquaintances. We take a template-based approach, and apply a classical Natural Language Generation (NLG) pipeline [8]:

Document planning. The first step is to analyze the requirements for the content of the text that has to be generated.

Our explanations included two main parts: **(1)** the reasoning behind the underlying mechanism of preference aggregation strategy; **(2)** the information of how group members' preferences and personalities played a role in generating the recommended song. Formulations for both of these parts are based on formulations from

previous work. For (1), we used an explanation template for the Additive Utilitarian aggregation strategy[1] [10]. *"Item X has been recommended to the group since it achieves the highest total rating"*.

We picked the part of explanation regarding how group members' preferences have been considered from Tran et al. [10] and the part for personality from Quijano-Sanchez et al. [7]. Below we use a working example for the scenario where the active user did not get their preference, but their acquaintances did.

1. Name: we picked parts of the template from [10] as follows: *"This decision supports the preferences of Bob and Carol .."*. We also add a negative component of the explanation, describing whose preferences have not been supported in this decision (e.g., *"This decision does not support the preferences of Ana .."*). Note: which parts are positive and which are negative depends on the scenario.

2. Rating: In a previous pilot study [6], we found that participants preferred to have categorization of preferences on a low-medium-high scale rather than as numeric ratings. To keep all explanations consistent and to reduce the number of variables for rating, we only considered *high* and *not high*. For example, in the scenario where the active user did not rate the song highly: *"The decision does not support the preferences of Ana who did not rate this song highly"*; and the others did prefer it: *"It supports the preferences of Bob and Carol who rated this song highly"*. Again, for whom the explanation uses 'highly' or 'not highly' depends on the scenario.

3. Personality: Inspired by Quijano-Sanchez et al. [7], we only show the personality of assertive members who have strong opinions and are difficult to convince. The member(s) with a strong opinion are always assumed to be the same users who got their preferred song. The scenario dictates whether this is the active user or their acquaintances. In our example, this was the acquaintances, so the explanation is: *"Besides, we have detected that Ana and Bob really want to listen to this song and won't be talked out of it easily."*

Discourse planning. The second step was to decide on the structure of the explanation. The structure was inspired by the *feedback sandwich model* [1]. The basic instruction for a feedback sandwich consists of one specific criticism (in our case the sentence about whose preferences has not been supported) "sandwiched" between two specific praises (in our case describing the mechanism and mentioning whose preferences have been supported).

Surface realization. To allow us to dynamically and automatically change the generated explanations we used the SimpleNLG[2] library for realizing natural language. This library helps handle combinations of parts of a sentence, punctuation etc. It also manages simple syntactic requests such as tense (e.g., past, present, future) and negation. After applying the aforementioned steps, we generate explanations such as the one in Figure 1.

3 EVALUATION AND DISCUSSION

We presented a framework which is adapted to users' privacy preferences to generate natural language explanations for groups.

Setup: To understand how much information the RS should expose to the group we asked users to adjust the explanations with the information they feel comfortable to share with their group

---

[1]This strategy takes into account the preferences of all individual group members. This explanation was found to be the most effective for user perceived fairness, consensus, and satisfaction.
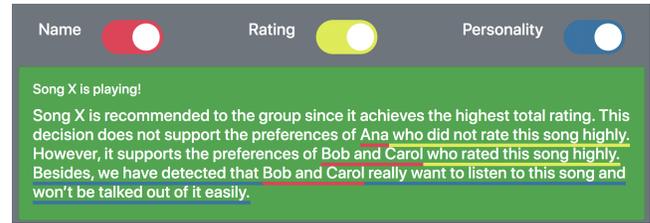[2]SimpleNLG (v. 4.4.8) is a "realisation engine", built by Albert Gatt and Ehud Reiter [2].

Figure 1: Screenshot of the system. Users can adjust the generated explanation using three different privacy controls: name, rating, personality. In this example, the explanation has all three controls enabled. The colors indicate the part of the explanation that each control influences. Here, Ana does not get her preferred song, but her acquaintances do.

members. The explanation components are all *off* by default, with participants having the option to turn them on. They were able to control three privacy-related option namely, whether to show/hide names, ratings, personality. The generated explanations are evaluated in a within-subjects user study (n=200).

Results: Percentages of the chosen privacy option to hide in the *Baseline*: name=32%, rating=**26%**, personality=**44%**; the *Unhappy user*: name=46%, rating=**30%**, personality=**48%**; the *Unhappy acquaintances*: name=**51%**, rating=**37%**, personality=47% [3]. We found that people use more privacy options in both low consensus scenarios compared to the Baseline, both differences were statistically significant (p <0.05, two sided McNemar-Bowker test).

Future work: We plan to conduct a user study with a live music recommendation setting, with real groups of various sizes. We also plan to extend the work to study the effect of privacy in other domains such as tourism.

REFERENCES

[1] Anne Dohrenwend. 2002. Serving up the feedback sandwich. *Family practice management* 9, 10 (2002), 43.
[2] Albert Gatt and Ehud Reiter. 2009. SimpleNLG: A realisation engine for practical applications. In *Proceedings of the 12th European Workshop on Natural Language Generation (ENLG 2009)*. 90–93.
[3] Judith Masthoff. 2011. Group recommender systems: Combining individual models. In *Recommender systems handbook*. Springer, 677–702.
[4] Shabnam Najafian and Nava Tintarev. 2018. Generating Consensus Explanations for Group Recommendations: an exploratory study. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*. ACM, 245–250.
[5] Vincent Robbemond Soumitri Vadali Shabnam Najafian Nava Tintarev Öykü Kapcak, Simone Spagnoli. 2018. TourExplain: A Crowdsourcing Pipeline for Generating Explanations for Groups of Tourists. In *ACM RecSys Workshop on Recommenders in Tourism*.
[6] Nivedita Prasad. 2018. *Evaluating Explanations for different Relationship Strengths*. Master's thesis.
[7] Lara Quijano-Sanchez, Christian Sauer, Juan A Recio-Garcia, and Belen Diaz-Agudo. 2017. Make it personal: a social explanation system applied to group recommendations. *Expert Systems with Applications* 76 (2017), 36–48.
[8] Ehud Reiter and Robert Dale. 2000. *Building natural language generation systems*. Cambridge university press.
[9] Nava Tintarev and Judith Masthoff. 2012. Evaluating the effectiveness of explanations for recommender systems. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 399–439.
[10] Thi Ngoc Trang Tran, Müslüm Atas, Alexander Felfernig, Viet Man Le, Ralph Samer, and Martin Stettinger. 2019. Towards Social Choice-based Explanations in Group Recommender Systems. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*. ACM, 13–21.

---

[3]The maximum and minimum frequencies are in bold.