Exploring User Concerns about Disclosing Location and Emotion Information in Group Recommendations

Najafian, Shabnam; Draws, Tim; Barile, Francesco; Tkalcic, Marko; Yang, Jie; Tintarev, Nava

DOI

[10.1145/3465336.3475104](https://doi.org/10.1145/3465336.3475104)

Publication date

2021

Document Version

Final published version

Published in

HT 2021 - Proceedings of the 32nd ACM Conference on Hypertext and Social Media

Citation (APA)

Najafian, S., Draws, T., Barile, F., Tkalcic, M., Yang, J., & Tintarev, N. (2021). Exploring User Concerns about Disclosing Location and Emotion Information in Group Recommendations. In *HT 2021 - Proceedings of the 32nd ACM Conference on Hypertext and Social Media* (pp. 155-164). (HT 2021 - Proceedings of the 32nd ACM Conference on Hypertext and Social Media). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3465336.3475104>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Exploring User Concerns about Disclosing Location and Emotion Information in Group Recommendations

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

Tim Draws
Delft University of Technology
Delft, the Netherlands
t.a.draws@tudelft.nl

Francesco Barile
University of Maastricht
Maastricht, the Netherlands
f.barile@maastrichtuniversity.nl

Marko Tkalcic
University of Primorska
Koper, Slovenia
marko.tkalcic@gmail.com

Jie Yang
Delft University of Technology
Delft, the Netherlands
j.yang-3@tudelft.nl

Nava Tintarev
University of Maastricht
Maastricht, the Netherlands
n.tintarev@maastrichtuniversity.nl

ABSTRACT

Recent research has shown that explanations serve as an important means to increase transparency in group recommendations while also increasing users' privacy concerns. However, it is currently unclear what personal and contextual factors affect users' privacy concerns about various types of personal information. This paper studies the effect of users' personality traits and preference scenarios—having a majority or minority preference—on their privacy concerns regarding location and emotion information. To create natural scenarios of group decision-making where users can control the amount of information disclosed, we develop *Toury-Bot*, a chat-bot agent that generates natural language explanations to help group members explain their arguments for suggestions to the group in the tourism domain. We conducted a user study in which we instructed 541 participants to convince the group to either visit or skip a recommended place. Our results show that users generally have a larger concern regarding the disclosure of emotion compared to location information. However, we found no evidence that personality traits or preference scenarios affect privacy concerns in our task. Further analyses revealed that task design (i.e., the pressure on users to convince the group) had an effect on participants' emotion-related privacy concerns. Our study also highlights the utility of providing users with the option of partial disclosure of personal information, which appeared to be popular among the participants.

CCS CONCEPTS

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

KEYWORDS

explanation; privacy concern; information privacy; group recommendation

This work is licensed under a Creative Commons Attribution International 4.0 License.

HT '21, August 30–September 2, 2021, Virtual Event, Ireland.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8551-0/21/08.

<https://doi.org/10.1145/3465336.3475104>

ACM Reference Format:

Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. 2021. Exploring User Concerns about Disclosing Location and Emotion Information in Group Recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media (HT '21)*, August 30–September 2, 2021, Virtual Event, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3465336.3475104>

1 INTRODUCTION

Explanations can be regarded as additional information that accompanies the recommendations and serves various goals, such as explaining the way the recommendation engine works to increase transparency [34]. Many studies have demonstrated the benefits of adding explanations to automated recommendations (e.g., [15, 33]). Previous research in this area has focused on explaining individual recommendations [15, 33]. When explaining recommendations to a group of users, it is challenging to recommend an item to a group that satisfies all group members simultaneously [1]. In particular, an additional aspect – users' *privacy* – has to be taken into account. In this context, although showing more information about group members could improve users' understanding of the recommendation process and perhaps make it easier to accept items they do not like, users' need for privacy is likely to conflict with their need for transparency [23], e.g., consider the explanation "*Alice is feeling sad today, and she really wants to visit this place*".

In our previous work [27], we found three factors that influence the trade-off between privacy and transparency in the tourism group recommendation context, namely, group members' *personality* (modeled using the Five-Factor Model [6]), *preference scenario* (whether the active user's preference is in the minority or majority compared to others' preferences within the group), and the type of *relationship* (the relationship strength between group members and equality of their positions). However, in that work, we looked at five kinds of different personal information (i.e., location, drug/alcohol, emotion, personal details, and personally identifiable information) shared in its entirety. So this needs further investigation to find out which of this personal information should be tailored for different personalities or group composition. Qualitative analysis from user comments indicated that for example people with different personality traits are concerned differently regarding different types of personal information [27]. The importance of the *information type* (i.e., the general category of the information that should be

disclosed) is also highlighted by Mehdy et al. [25]. In this work, we consider the tourism domain and consider the specific personal information types individually rather than in their entirety, namely location and emotion, which are most used in the current tourism recommender systems (e.g., [26]). Another main distinction to our previous work is that we looked at participants' privacy decisions through their actual behavior rather than only their attitude, which is the case in most privacy-related research. Specifically, we investigate the following research questions:

RQ1: How do personality and preference scenario affect people's location-related privacy concerns in explaining group recommendations to their group?

RQ2: How do personality and preference scenario affect people's emotion-related privacy concerns in explaining group recommendations to their group?

To answer the above research questions, we designed a user study where participants receive recommended *point of interests* (POIs) from the group recommender in both majority and minority preference scenarios. Depending on the preference scenario, they are instructed to convince the group either to visit or skip a recommended place, by explaining their arguments for suggestions to the group. To facilitate the study, we developed *TouryBot*, a chat-bot agent that supports the natural dynamics of group decision-making. Due to the diverse needs and preferences, recommendations for *groups* are particularly challenging that often require discussions among group members. Our chat-bot allows us to control the flow of information by suggesting gradual revealing of information to users; at the same time, it improves the ecological validity of people chatting together about potential POIs.

Our results indicate that users generally have a larger concern regarding the disclosure of emotion compared to location information. In contrast to previous research, we find no evidence that personality traits or preference scenarios affect privacy concerns. Further analysis revealed that task design had a strong effect on participants' emotion-related privacy concerns. In particular, the *nudging* of users to convince the group can partly explain such lack of evidence. Our study also reveals the utility of providing users with the option of partial disclosure of personal information, which appeared to be popular among the participants to strike a balance between transparency and privacy. Our results, therefore, in addition to discussing our research questions, show the effects of relevant design choices – i.e., nudging people to convince the group and providing the option for partial information disclosure – that should be taken into account when designing chat-bots and similar tools for decision-making in group recommendations contexts.

All material for analyzing our results and replicating our user study, (i.e., chat-bot implementation, user study materials, data gathered in the user study, and the analysis scripts) is publicly available: <https://osf.io/6bfpd>.

The remainder of this paper is structured as follows. Section 2 introduces related work in the area of privacy concern in explanations for group recommender systems and group deliberation, which arise in these scenarios. Section 3 presents the user study performed to investigate the privacy aspects of explaining to groups. Section 4 presents the results and analysis of our user study, while Section 5 discusses the main findings and presents the limitations

of our approach, and provides recommendations for future work. Finally, Section 6 summarizes our findings.

2 BACKGROUND AND RELATED WORK

This section provides an overview of existing research related to explanations in group recommender systems and discusses privacy aspects in explanations for groups, which arise in these scenarios, focusing on findings related to our research question. We also discuss related work on interfaces and systems to support interactions in group decision-making processes, which inspire the choice and design of chatbot agents in our study.

2.1 Group Privacy Concern

Existing works on explanations for recommendations mostly focus on the benefit of transparency, i.e., increasing users' understanding of the system's reasoning in recommendation generation [34]. However, when generating explanations for groups rather than individuals, privacy becomes of great relevance. Najafian et al. [28] investigated which information people would like to disclose in explanations for group recommendations in the music domain. The work was recently extended [27] to evaluate the factors that have an impact on privacy concerns for group recommendations. They show an impact deriving from (i) the personality, (ii) the preference scenario and (iii) the relationship type. Furthermore, Mehdy et al. [25] suggested to consider the information type when modeling users' situation-specific privacy concerns. In the following subsections, we discuss relevant literature on these factors and formulate the hypotheses that lead this work.

2.1.1 Personality. Several studies in the field of behavioral sciences analyze the impact of personality on an individual's privacy concerns. The results, however, are not consistent with each other. Personality is generally modeled using the Five Factors Model (FFM), also known as Big Five or OCEAN. It models individuals' personality with five traits: Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism [6]. Bansal et al. [3] analyzed the effect of the individual personality on information disclosure in three classes of websites (Finance, E-commerce, and Health). Their results showed a significant positive effect of both Agreeableness and Neuroticism. In the context of location-based services, Junglas et al. [20] showed significant effects of Agreeableness but suggesting a negative effect (i.e., more agreeable people were less concerned about their privacy). In the context of explanations for group recommendations, Najafian et al. [27] showed that more agreeable and extroverted people were more concerned with privacy. Related work shows mainly that the three personality traits *agreeableness*, *extraversion*, and *neuroticism* are related to privacy concern [3, 20, 22].

2.1.2 Preference scenario. Several studies suggest that the preference scenario within the group could have an impact on privacy concerns. In particular, people having minority preferences compared to others' preferences within the group could decide to not share their preferences in order to match the opinions of the majority, for a phenomenon known as conformity [2, 10]. This was

confirmed by Najafian et al. [27], which showed that people having minority preferences expressed higher privacy concerns, in particular for the information related to their emotions.

2.1.3 Type of the Relationship. Social relationships have been shown to be a contextual factor that has an impact on privacy concerns in information sharing [9, 12, 16, 25]. For example, it has been found that people have a more positive attitude toward information disclosure to recipients with a close relationship (i.e., a family member, or a friend) than to those with a weak relationship (i.e., colleagues) [25]. Additionally, Wang et al. [36] proposed adding to the strength of the relationships (which they call tightly versus loosely coupled) a second dimension considering the relative standing or position within the group: (i) positionally homogeneous (i.e., groups where the position of the members are equal, as a group of friends) and ii) positionally heterogeneous groups (in which the position is unequal, as a family). Following this classification, Najafian et al. [27] showed that privacy concerns are perceived more in loosely-coupled heterogeneous groups than tightly-coupled homogeneous. In this work, we therefore focus on loosely coupled (weak ties) heterogeneous groups to consider privacy concern in an extreme case.

2.1.4 Information type. Existing studies have suggested that the magnitude of privacy concerns depends on the type of information to disclose. In this work, we focus on privacy concern in a group tourism scenario. Previous work in this area proposed a context-aware recommender system for tourism that use users' current *location* and *emotion* (mood) to generate personalized recommendations [26], while Najafian et al. [27] highlighted location and emotion information as the information that generates the higher privacy concerns in the context of group recommendations in tourism. Additionally, Tsai et al. [35] highlighted how allowing users to control the granularity level of the shared location information could decrease the related privacy concerns, although this can reduce the benefits of sharing the information in several application scenarios. Finally, Consolvo et al. [5] highlight that the level of detail of the requested information is important, as the users are willing to just disclose the amount of information they think are useful according to the specific scenario or deny the request.

Since previous work has found an effect of personality and preference scenario on privacy concern (c.f., Section 2.1.1 and Section 2.1.2), we study these here specifically in a group context. Specifically, we formulate the following hypotheses regarding location-related privacy concern (measured as the level of location information disclosed):

- H1.a)** Extraversion affects location-related privacy concern.
- H1.b)** Neuroticism affects location-related privacy concern.
- H1.c)** Agreeableness affects location-related privacy concern.
- H1.d)** People with minority preferences have a higher location-related privacy concern than people with majority preferences.

Caliskan Islam et al. [4] categorized emotion as private information. Graham et al. [11] showed that the expression of negative emotion is useful to elicit help from others and that people who are more willing to express negative emotion have larger social networks. They underlined, however, the need of expressing such

emotions in a way that is appropriate to the particular situation and with people with whom a relationship has been established.

Led by these findings, we formulate the following hypotheses on the impact of personality and preference scenario on the privacy concerns related to emotion-based privacy concern (measured as level of emotional information disclosed):

- H2.a)** Extraversion affects emotion-related privacy concern.
- H2.b)** Neuroticism affects emotion-related privacy concern.
- H2.c)** Agreeableness affects emotion-related privacy concern.
- H2.d)** People with minority preferences have a higher emotion-related privacy concern than people with majority preferences.

2.2 Group Deliberation

Presenting users with a static recommendation list does not consider scenarios in which the user might construct their preferences during the decision-making process [18]. This is especially true in scenarios where the target users are not individuals but a group of people. In such cases, the group choice depends not only on individual preferences at the beginning but also on the dynamics of group discussion when making joint decisions [30]. Therefore, previous research has introduced strategies to enable interaction between group members during the process [17, 24].

Conversational interfaces specifically have shown to lead to higher satisfaction of the user, requiring less interaction and increasing the likelihood of them using the system in the future [13]. In this direction, Nguyen and Ricci [31] presents a system that allows the group members to revise their preferences through a conversational process with a chat-based interface, showing how that can increase the system usability and the recommendation quality. So in contrast with their system that suggests individual recommendation to each member, in our system, the chat-bot suggests what is best for the group to each group member and it is up to that person to share it with the group. Then our chat-based system offers the possibility to support the decision-making process by providing natural language explanations of the recommendations given, or by supporting users in a group discussion by suggesting arguments for their positions.

3 METHOD

In this section, we describe an online between-subjects study that investigates how personality and preference scenario relate to individuals' privacy concern about disclosing their current location (e.g., "John is on Vondelstraat (a central station in Amsterdam) and emotion information (e.g., "John is feeling grief") in a group recommendation explanation.

3.1 Study Platform

To answer the research question we implemented a web-based chatbot that we call *Tourybot*. For the UI we used a client in java (Vaadin AI Chat) ¹ and implemented in Vaadin framework.² The backend is written in python. SQLite was used for logging user interactions in the task. Tourybot includes two chat windows, one for the chat

¹<https://github.com/alejandro-du/vaadin-ai-chat>, retrieved March 2021.

²An open platform for building web apps in Java (<https://vaadin.com/>), retrieved March 2021.

between the system bot and individual members (see Figure 2), and the other for the chat with the Group (see Figure 1). Users can seamlessly switch between the two chats to add system-generated recommendations and explanations to their discussions with other group members.

3.2 Measures

Inspired by previous work [27], our study considers an experimental manipulation of users' preference scenarios; i.e., either having *minority preferences* or *majority preferences* in the group. The relationship type among group members was (in both cases) predefined as a "loosely coupled heterogeneous group" (e.g., a lecturer and students). We additionally observed users' personality traits and included location-related, emotion-related privacy concerns as dependent variables.

3.2.1 Independent variables.

Preference scenario (binary). Each participant in our study was exposed to either *minority* or *majority* preference scenarios, tasked to convince the group to either skip or visit a POI through explanations that are privacy-sensitive.

- **Minority preferences:** the active user's preference is in the minority within the group. An item that is not the (active) user's favorite has been suggested to the group by other (synthetic) group members. In this case, the participant tries to convince others to skip the recommended POI. This creates a trade-off between disclosing more personal information (risking privacy violation) and going to a POI they are not interested in.
- **Majority preferences:** the active user's preference is in the majority within the group. An item that is the user's favorite has been suggested to the group. In this case, the participant tries to convince others to visit the POI. This creates a trade-off between disclosing more personal information and missing a POI they want to visit.

Personality (continuous). We used the *Big Five Inventory* (BFI) to assess individuals' personality on the three traits of *Extraversion*, *Agreeableness*, and *Neuroticism* [19]. The questionnaire is composed of 44 questions with a 5-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree). Responses are aggregated by taking their mean.

3.2.2 Dependent variables.

Location-related privacy concern (ordinal). We used three different levels of granularity for location-related information to measure the group members' privacy concern regarding that information being disclosed in the group. Users had three options to choose from: "*no location*" (value of 1) as has been considered as low-level granularity or not sensitive, "*neighborhood location*" (value of 2) as has been considered as middle-level granularity or medium sensitive, and "*exact location*" (value of 3) as has been considered as high-level granularity or very sensitive.

Emotion-related privacy concern (ordinal). Similarly, we used three different levels of granularity for emotion-related information to measure the group members' privacy concern regarding

that information being disclosed in the group: "*no emotion*" (value 1) as has been considered as low-level granularity or not sensitive, "*mild emotion*" (value 2) as has been considered as middle-level granularity or medium sensitive, and "*intense emotion*" (value 3) as has been considered as high-level granularity or very sensitive.

3.2.3 Descriptive measures. We collected participants' age and self-identified gender to enable a demographic description of our sample. Participants also stated how familiar they are with the city in recommendation (Amsterdam) by responding on a 5-point Likert scale (ranging from "not at all familiar" to "extremely familiar"). However, familiarity with Amsterdam did not make any difference on the results.

3.3 Materials

3.3.1 Emotion content. This study concerns users' willingness to disclose emotion- and location-related information, which are among the five personal information types used in previous work [27]. These two information types have been used in current tourism recommended systems; e.g., Mohamed et al. [26] use users' current location and emotion (or mood) to recommend personalized travel-related POIs to visit. We conducted a pre-study to identify which specific emotion would best lend itself to be included in the scenario we would present to participants in the main study. To do this, we aimed to verify which emotion-related information could raise privacy concerns in participants to be included in the explanation. Note that no pre-study was conducted for location-related information as privacy concern about disclosing current location has been studied extensively (e.g., [35]).

Ekman and Friesen [8] identify six basic emotions. Each of them has a corresponding *intense form* (i.e., rage as intense form of anger, loathing as intense form of disgust, terror as intense form of fear, ecstasy as intense form of happiness, grief as intense form of sadness, and amazement as intense form of surprise) [32].

To decide which emotion to include in the study, we asked 18 students at our university to imagine planning an activity with a group of people that they don't feel very close with, using a group chat. Furthermore, the social positions of the group members are not equal. For example, the group could consist of a lecturer and some students, where the participant is one of the students (i.e., a loosely coupled heterogeneous). They were asked how comfortable they would be, in such a scenario, in sharing their emotions in the group chat to explain and support their arguments, by responding on a 5-point Likert scale ranging "extremely uncomfortable" to "extremely comfortable". We asked students to perform this evaluation for each of the six basic emotions and their corresponding intense form. We also allowed participants to indicate additional emotions which they considered sensitive to disclose.

We conducted a Repeated Measures ANOVA to analyze whether participants in the pre-study had different levels of comfort regarding the disclosure of the different emotions. Indeed, we found a significant difference ($F = 19.57, p < 0.001$). Among the different emotions, participants were on average least comfortable with sharing *sadness* (mean = 2.06, sd = 0.94) and its corresponding intense form *grief* (mean = 1.67, sd = 0.84). We thus chose this combination of sadness (basic emotion) and grief (intense emotion) for our study.

Figure 1: The chat in the majority preference scenario between an active user and his group. a) indicates two ongoing chats, one with a chat-bot and the group, b) indicates the active user (John) shares his preference (the Van Gogh museum in this example) with his two group members (the two other group members, Bob and Alice, are hypothetical) in a group chat.

Figure 2: The chat in the majority preference scenario between the chat-bot and an active user. a) indicates two ongoing chats, one with a chat-bot and the group, b) indicates an ongoing chat with a chat-bot. Here the user can indicate the level of location information they want to share to convince the other group member (Bob) to visit the suggested POI.

3.3.2 Initial POIs. For the user study, we needed POIs for both the minority and majority preference scenarios. To collect such POIs, we provided participants with ten initial POIs to rate on a 5-point Likert scale (ranging from “definitely would not visit” to “definitely would visit”). The ten initial POIs retrieved from the most frequently visited POIs in the city of Amsterdam from the social location service *Foursquare*.³ By using participants’ real preferences, we aimed to increase the likelihood of a more realistic situation for users to imagine.

³<https://developer.foursquare.com/>, retrieved February 2021.

3.4 Procedure

Participants who accepted our task received brief instructions about the task and were asked to check off an informed consent before beginning their task session. After consent for the study participants went through the following steps.

Demographics & Preferences (Figure 3a). Participants first completed a short demographic questionnaire. They were also asked their first name and to form their (hypothetical) group by naming two people of whom they thought that 1) their social positions were

Figure 3: Overview of the experiment procedure for each participant: the system presents a) a demographics questionnaire and the preference elicitation step, b) either minority or majority preferences scenario which includes questions both about location and emotional information, c) the final personality questionnaire. Arrows indicate the order of information types are randomized.

unequal and 2) they are not close to each other, i.e., communication among them is not frequent (e.g., the group could consist of a lecturer and some students, where the participant is one of the students). We also elicited their preferences as described in Section 3.3. Note that the group always consisted of three group members, with only one of them being an active user and the other two being hypothetical group members.

Preference Scenario (Figure 3b). Participants were randomly assigned to take part either in minority or majority preferences scenario. If they were assigned to the minority preferences scenario, they were asked how much information regarding either their location or emotion they are okay to share for an imaginary POI in Amsterdam with their hypothetical group members (the order of information types was also randomized). We informed them that these were not their real information, but they should imagine that it is correct also not shared with anyone external. They had three options for how much they are okay to share with their group as explained in Section 3.2.2. The active user could use those options to provide more information about his current location or emotion to support their argument to skip the suggested place by the hypothetical group member.

We toggled between skipping a POI (having minority preferences) and visiting a POI (having majority preferences) in the way we convince the group. If they were assigned to the majority preferences scenario, the active user tried to convince the group to visit to the suggested POI by the user by providing more information to support the arguments. In Figure 1, Bob (the hypothetical group member) asks why John (the active user) suggests that POI. And as can be seen in Figure 2, the active user tried to provide more information about his current location to support his arguments.

To terminate the dialog with the chat-bot in either scenario, we asked users if they are okay now with the current explanation to the group or whether they still wish to edit it. Participants had the option to go back to the information they chose already to disclose more information to their group, as in real situations they could

not decrease the amount of information they already shared in the group. For the analysis, we only considered users' final decisions.

The two information types were randomized between participants, to prevent biases due to ordering or learning effects. The options for how much information to show to the group were ordered based on the information hierarchy from low information to high information for example from no location to exact location.

To ensure that users read all the relevant conversations, we did two things: 1) showing a pop-up to the user to switch to the other window when needed, and 2) duplicating the messages to make sure that the user does not miss any information. For example, when in group chat, a simulated conversation by a hypothetical group member (Bob) asked "Why are you suggesting to go here?", we showed a pop-up saying "you have a new message in the Tourybot chat" and a button to switch to the TouryBot chat. In the TouryBot chat, we repeated this message at the beginning of the conversation with the active user (as can be seen in Figure 2b).

Personality (Figure 3c). After completing the scenario, participants filled in the BFI for assessment of their personality traits.

3.5 Participants

To determine the required sample size, we performed a power analysis [7] for a between-subjects experiment. Assuming medium effects for all four factors (i.e., preference scenario and three personality traits; odds ratio = 3) and otherwise assuming that participants who are (a) in the majority setting and (b) have medium levels across the personality scales are equally likely to choose between the three location or emotion preferences, we arrived at a recommended sample size of 360. We recruited 374 participants from the crowdsourcing platform Prolific.⁴ This platform has shown to be an effective and reliable choice for running relatively complex and time-consuming studies, e.g., for interactive information retrieval [37]. To ensure reliable participation, we followed Prolific guidelines and restricted eligibility to workers who had an acceptance rate of at least 80% and had at least 10 successful submissions on the platform. We paid participants the wage suggested by Prolific. We excluded from our results participants who failed at least one attention check.

The resulting sample of 362 participants had an average age of 33.4 (sd = 13.5) with a satisfactorily balanced gender distribution (51% female, 38% male, 11% other – which also includes those who did not answer to this question).

3.6 Statistical Analyses

To test our hypotheses (see Section 2), we performed two *ordinal logistic regression* (OLR) analyses [14] (i.e., one to predict privacy concerns regarding location-related information and the other for emotion-related information) with *preference scenarios* and the five personality scales *extraversion*, *agreeableness*, *neuroticism*, *openness*, and *conscientiousness* as independent variables.⁵ We corrected for multiple hypothesis testing by lowering the significance threshold to $\frac{0.05}{8} = 0.00625$ (i.e., applying a *Bonferroni correction* [29]).

⁴<https://www.prolific.co>

⁵ Although our hypotheses concerned only the first three of the five personality scales (see Section 2), we added *openness* and *conscientiousness* as covariates to the models to account for potential confounding factors.

4 RESULTS

In this section, we discuss the outcomes of the hypothesis tests we conducted and present several exploratory findings.

4.1 Hypothesis Tests

Table 1 shows the results from the OLR analyses regarding location information and emotion information. We found *no evidence* in favor of any of our eight hypotheses (**H1a** - **H2d**; all $p > 0.00625$; see also Section 3.6). In contrast, the *odds ratios* (OR) of the regression factors we tested suggest that users were approximately equally likely to have higher location-related privacy concern (i.e., disclosing their exact location, their neighborhood location, or no location), holding constant all other variables, across different levels of *extraversion* (OR = 0.94, 95% CI[0.70, 1.25]; **H1a**), *agreeableness* (OR = 0.81, 95% CI[0.56, 1.17]; **H1b**), and *neuroticism* (OR = 1.06, 95% CI[0.80, 1.40]; **H1c**), as well as different preference scenarios (i.e., minority and majority preferences in the group; OR = 0.91, 95% CI[0.60, 1.38]; **H1d**). Similarly, users were approximately equally likely to have higher emotion-related privacy concern (i.e., disclosing their exact emotion, their approximate emotion, or no emotion), holding constant all other variables, across different levels of *extraversion* (OR = 0.83, 95% CI[0.63, 1.09]; **H2a**), *agreeableness* (OR = 0.93, 95% CI[0.65, 1.31]; **H2b**), and *neuroticism* (OR = 0.87, 95% CI[0.67, 1.14]; **H2c**), as well as different preference scenarios (OR = 0.66, 95% CI[0.45, 1.97]; **H2d**).

In sum, based on the OLR results, we cannot reject any of the null hypotheses opposing the hypotheses we aimed to test (see Section 3). Odds-ratios computed as part of these analyses suggest that the hypothesized effects (i.e., of the three personality traits and preference scenario on location-related and emotion-related privacy concern) may be absent or much smaller than previously anticipated in this context.

4.2 Exploratory Findings

In this section, we present several exploratory findings that may help to explain the results from the hypothesis tests.

4.2.1 Familiarity. Most participants were not familiar with the city in recommendation (Amsterdam), as 85% of them selected one of the bottom three options from the Likert scale. Moreover, familiarity was unrelated to location- or emotion-related privacy concern ($p = [0.47, 0.90]$; results of ordinal logistic regressions).

4.2.2 Partial Disclosure. Table 2 shows that nearly half (40%) of participants chose to *partially* disclose both location-related and emotion-related information (i.e., disclosing their neighborhood location or approximate emotion) rather than fully hiding or disclosing it.

4.2.3 Information Type. In line with previous research [25], we found that privacy concern varies depending on information types, with significantly larger concern for emotion-related compared to location-related information ($V = 4831.5$, $p < 0.001$; result of a Wilcoxon signed rank test with continuity correction). Table 2 shows that, whereas 33% (122) of participants did not share any emotional information, only 5% (19) of participants chose not to share any location information with their group.

4.2.4 Task Completion Time. Participants who were exposed to the minority preferences scenario spent more time performing the task (mean = 132.1s, sd = 65.5s) compared to participants who were exposed to the majority preferences scenario (mean = 103.5s, sd = 53.6s; $t = 4.55$, $p < 0.001$; result of an independent samples t -test). This shows that, although participants disclosed similar amounts of information in the two preference scenarios, they may be more hesitant in doing so when placed in the minority. This might be because, in this scenario, they had a more difficult time to give away some information to convince other group members to skip the suggested place. In line with this, we found that 70% of participants who changed their privacy setting to disclose more information at the final step of the study were participants in the minority preferences scenario.

4.2.5 Task Design. Our study exposed users to one of two tasks: either (1) to convince other group members to *accept* visiting the suggested POI or (2) to convince other group members to *skip* the suggested place. Both tasks thus required participants to *convince* other group members. Therefore, one explanation for why we did not obtain the expected results is that our task design nudged participants into a “convincing mindset”. This could have caused the effects of personality traits and preference scenarios – which have been demonstrated in previous studies – to decrease to a degree that we could not pick them up in this study. To investigate if the task design of convincing other group members had an effect on our study regarding privacy concern, we adapted the task to a scenario in which the active user was still placed in either a minority or majority preference but where the other hypothetical group member did not push asking questions regarding the active user suggestion. Instead, the hypothetical group member in this new design would simply agree with the active user even before any location-related or emotion-related information was disclosed. The only change in this new design compared to the task described in Section 3.4, thus, were the questions asked by the hypothetical group member which adapted to i.e., for minority scenario: “*that’s alright, we can skip this place*”, and for the majority scenario: “*that’s alright, we can visit this place*”.

We recruited an additional 200 participants through *Prolific* with the exact same conditions as the first study (see Section 3.5).⁶ 179 participants remained after removing those who failed the attention checks. We added this additionally obtained data to our original data set of 362 participants, resulting in a data set containing 541 observations (i.e., 362 of which came from the convincing task design and 179 of which came from the non-convincing task design). This allowed us to run the OLR analyses again with *convincing* as an additional factor. Otherwise, the OLR analyses were performed in the same way as described in Section 3.6.

Whereas *convincing* did not have an effect regarding location-related privacy concern ($\beta = -0.25$, $p = 0.17$, OR = 0.78 with 95%CI[0.55, 1.11]), it did affect emotion-related privacy concern ($\beta = -0.97$, $p < 0.001$, OR = 0.38 with 95%CI[0.26, 0.54]). This means that, when people had to *convince* other group members in our first task design, they disclosed more emotion-related information compared to our second task design where they didn’t

⁶A required sample size of 180 additional participants was computed in a simulation study beforehand.

Table 1: Results from two OLR analyses with location-related privacy concern (left) and emotion-related privacy concern (right) as dependent variables (DVs). Factors included two intercepts (i.e., due to the three-level, ordinal dependent variables), preference scenario (pref) and the five different personality scales extraversion (extr), agreeableness (agr), neuroticism (neur), openness (open), and conscientiousness (cons). Per factor, we report the β regression coefficient, p -value, and OddsRatio (OR; with 95% confidence interval; CI). We tested some of these factors as part of our hypothesis tests (see Section 2). However, no factors were statistically significant after correcting for multiple testing (see Section 4).

DV: Location-Related Privacy Concern					DV: Emotion-Related Privacy Concern				
Hyp.	Factor	β	p	OR [95% CI]	Hyp.	Factor	β	p	OR [95% CI]
-	Intercept 1 2	-1.11	0.38		-	Intercept 1 2	-1.83	0.12	
-	Intercept 2 3	1.66	0.19		-	Intercept 2 3	-0.10	0.93	
H1a	extr	-0.07	0.65	0.94[0.70, 1.25]	H2a	extr	-0.19	0.18	0.83[0.63, 1.09]
H1b	agr	-0.21	0.26	0.81[0.56, 1.17]	H2b	agr	-0.08	0.66	0.93[0.65, 1.31]
H1c	neur	0.05	0.71	1.06[0.80, 1.40]	H2c	neur	-0.14	0.31	0.87[0.67, 1.14]
-	open	-0.17	0.31	0.84[0.60, 1.18]	-	open	-0.11	0.51	0.90[0.65, 1.24]
-	cons	0.04	0.78	1.04[0.77, 1.42]	-	cons	0.27	0.06	1.32[0.99, 1.75]
H1d	pref	-0.09	0.66	0.91[0.60, 1.38]	H2d	pref	-0.42	0.03	0.66[0.45, 0.97]

Table 2: Number (and percentage) of participants across privacy concerns regarding location (left) and emotion (right).

Exact Loc.	Neighborhood Loc.	No Loc.	Exact Emot.	Approximate Emot.	No Emot.
202 (55%)	149 (40%)	19 (5%)	100 (27%)	148 (40%)	122 (33%)

have to convince other group members. Although we did not find the same effect regarding location-related privacy concern, the trend there went in the same direction as for emotion-related privacy concern. It could thus be that the convincing aspect affected location-related privacy concern to a lesser extent and that we did not collect enough additional data to pick it up. This would be in line with the exploratory findings reported in Section 4.2.3 (i.e., that people are generally more willing to disclose location-related information compared to emotion-related information). Finally, it should be pointed out that, although we found that participants disclosed less emotion-related information in the non-convincing task design, they still did not differ across personality traits or preference scenarios in this adapted context. This could be due to additional confounding factors that we did not measure here.

4.3 Qualitative feedback

Partial disclosure. In line with the results, people seem to be happy to have the partial disclosure option to balance between their need to convince other group members and their need to not violate their own privacy.

For example, “I liked that there was the option to share approx location rather than exact”, or, “I didn’t want to give too much information away to people I didn’t know, but I wanted to be able to give good enough reasons for my choices.”, and another one, “I wanted to share my approximate location and approximate emotion to try to convince ... that going to the veggie restaurant was a good idea but I did not want to go into too much specific detail about how I was feeling because we do not know each other well and that felt too personal to share in a group chat.”

Changing their mind. Only a few participants (16%) changed their first selected options of disclosing information and actually disclose more with their groups. Interestingly 69% of those were high neurotic people and mainly they mentioned they nudged to disclose more information.

For example, “... it felt embarrassing to provide the exact emotion, but the group members were argumentative and kept pushing, so it felt like I needed to justify myself” or “I offered less information at first then added more in an effort to convince the other members.”, or, “i don’t like revealing information about myself unless it is necessary.”

Relationship. In this study, we kept the relationship constant, however, 10% of participants explicitly mentioned the effects of the relationship that caused them to share less.

For example, “I would be happy to offer my opinion on where I would like to go in Amsterdam, but I would not be comfortable sharing my emotions with people I do not know very well. If I was in Amsterdam with close friends I would tell them I am feeling depressed.”

Chat-bots design. Overall, more than 80% of participants greatly enjoyed using the chat-bots and found it unique, engaging, interactive, suitable for planning their trip with a group, and potential for actual products.

For example: “I enjoyed filling out the study. seems a good idea for planning a trip.”, or, “It was a cool, interactive and interesting study, much more interesting than many others.”, or, “study was really engaging and different, I had a really good time taking part in it.”, or, “Think the study was interesting and has potential for actual products.”

5 DISCUSSION

We presented a user study to investigate the effects of three personality traits (i.e., extraversion, agreeableness, and neuroticism) as well as preference scenario (i.e., having minority or majority preferences) on users' privacy concern in a realistic chat-bot scenario (see Section 3). Our results contain no evidence for any of these effects (Section 4). In contrast, the odds ratios we computed suggest that the effects we aimed to investigate may not be present in the context of our study.

These findings are not in line with previous research that suggests that the preference scenario, as well as personality traits, do affect users' privacy concern [27]. In this section, we discuss our results. We describe several potential reasons for why they contrast previous research in this area and highlight important implications for the design of chat bots and similar tools that aim to bridge the gap between group recommendation systems and consumers.

Task Design. Our task design had an effect on emotion-related privacy concern (see Section 4), which could be one of the reasons for not finding previously demonstrated effects in this set-up. This also implies that it is important for such chat bots to avoid nudging people into some "convincing" mindset, as they might disclose more information than they are comfortable with.

POI Sensitivity. To make the scenario more realistic for participants, in this study we used regular POIs; i.e., 10 most frequently visited POIs from Foursquare's five main categories (e.g., *Arts & Entertainment and Food*). This was different compared to previous research that reported effects of preference scenario and personality traits on privacy concern, where particularly a sensitive POI was used (e.g., a cannabis store) [27]. The arguably lower overall privacy sensitivity in our study might have diminished these effects, e.g., causing people having minority preferences to feel less concerned regarding the disclosure of their personal information compared to these previous studies.

Partial Disclosure. To provide users with an easier option (to be able to give away some part of their information to convince other group members but still not disclose all their personal information) rather than only disclose or hide their personal information, in addition to those extreme options, we offered partial disclosure of this information as well. The high number of selections of this option for both information types (40%), suggests this can be a beneficial option to offer in such a group explanation context. Besides, we argue adding this middle option to this study rather than the previous studies which only provided a show and hide options might cause some participants who would normally choose either of the extreme options (show/hide) to choose this middle option.

Information Types. The two types of information we included in this study were location and emotional information. This decision was based on previous results that showed that people had concern regarding disclosing these types of information in a group explanation [27]. Furthermore, in our study especially 40% chose partial disclosure of each information type (e.g., neighborhood location, approximate emotion) which shows people do have privacy concern regarding these information overall.

Task Completion Time. We found significant task completion time differences between participants exposed to the minority preferences scenario with those who were exposed to the majority preferences scenario. This suggests that participants may have had difficulty deciding but then went with it.

Implications. Based on the discussed results we argue providing a partial disclosure option in designing chat-bots would make it easier for users to balance their need for transparency while not violating their privacy by disclosing too much personal information. On the other hand, our study design diminished previously demonstrated effects, which might be important for designing chat-bots in such a group recommendation scenario to avoid nudging people into some "convincing" mindset as in this context they might disclose more personal information than they are comfortable with.

5.1 Limitations and Future Work

General Privacy Concern. In this study, we did not find an effect of personality traits and preference scenario on privacy concerns. There might be an additional mediating factor that affects participants' privacy concerns. For example, it would be beneficial in future work to also measure general privacy concern [21] to see if it mediates privacy concern regarding disclosing their personal information in a group.

Hypothetical Group. In this study, we only had one active user to control the group scenario in a way to see if there will a group member who needs to be convinced how much information the active user is OK to share in the group. For future work, it would be more realistic to use real groups to see how group members deal with this tension of disclosing more information to convince other users to accept what they want but on the other hand not violating their privacy by disclosing too much personal information.

Constant Relationship. In this study, we only picked one type of relationship that in the previous study has been shown to have higher privacy concerns. However as some participants (10%) explicitly mentioned in their comments, that the type of relationship affected their choice to share less emotion or location information with their group, this is an interesting future work to study the effect of the relationship on privacy concern in more details.

User Control. To be able to study participants' privacy concern we only provide options that they can argue their choices with other group members (or only provide more information for their suggestions in case of the new, not convincing task design). However, as stated in the comments as well mainly high agreeable participants needed an option to just accept what other group members suggest and did not want to argue with them. In a future study, all people to decide themselves if they convince other group members or if they want to go along with other group members' suggestions.

6 CONCLUSIONS

We presented a user study investigating the effect of three personality traits and preference scenarios on user's privacy concerns. In line with the previous study, we found that privacy concern varies depending on information types, with significantly larger concern for emotional than location information. However, surprisingly in contrast to the previous studies, neither the three personality traits that we measured nor the preference scenario affected this privacy concern. One explanation for why we did not obtain the expected results could have been that our task design nudged participants into a "convincing mindset". Therefore, we recruited additional participants to study the effects of task design, and we found that it had an effect on participants' emotion-related privacy concern but not location-related privacy concern. These results suggest that it is important for chat-bots to avoid nudging people into a "convincing" mindset as in this context they might disclose more personal information than they are comfortable with. Additionally, we found that users used partial disclosure very often. It seemed to make it easier for users to handle the trade-off between their need for transparency to convince other group members and their need for preserving their privacy by not disclosing too much personal information. Another recommendation for the design of chat-bots which support decision-making in group recommender systems context, is therefore to provide users with an option for partial disclosure (in addition to the complete and no disclosure options).

ACKNOWLEDGMENTS

This publication is partly financed by the Dutch Research Council (NWO). This work made use of the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-1776.

REFERENCES

- [1] Kenneth J Arrow. 1950. A difficulty in the concept of social welfare. *Journal of political economy* 58, 4 (1950), 328–346.
- [2] Solomon E Asch. 1956. Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological monographs: General and applied* 70, 9 (1956), 1.
- [3] Gaurav Bansal, Fatemeh Mariam Zahedi, and David Gefen. 2016. Do context and personality matter? Trust and privacy concerns in disclosing private information online. *Information & Management* 53, 1 (2016), 1–21.
- [4] Aylin Caliskan Islam, Jonathan Walsh, and Rachel Greenstadt. 2014. Privacy detective: Detecting private information and collective privacy behavior in a large social network. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, 35–46.
- [5] Sunny Consolvo, Ian E Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. 2005. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 81–90.
- [6] Paul T Costa and Robert R McCrae. 1992. *Neo personality inventory-revised (NEO PI-R)*. Psychological Assessment Resources Odessa, FL.
- [7] Patrick Dattalo. 2008. *Determining sample size: Balancing power, precision, and practicality*. oxford university press.
- [8] Paul Ekman and Wallace V Friesen. 1986. A new pan-cultural facial expression of emotion. *Motivation and emotion* 10, 2 (1986), 159–168.
- [9] M Max Evans, Ilja Frissen, and Chun Wei Choo. 2019. The Strength of Trust Over Ties: Investigating the Relationships between Trustworthiness and Tie-Strength in Effective Knowledge Sharing. *Electronic Journal of Knowledge Management* 17, 1 (2019).
- [10] Donelson R Forsyth. 2018. *Group dynamics*. Cengage Learning.
- [11] Steven M Graham, Julie Y Huang, Margaret S Clark, and Vicki S Helgeson. 2008. The positives of negative emotions: Willingness to express negative emotions promotes relationships. *Personality and Social Psychology Bulletin* 34, 3 (2008), 394–406.
- [12] Mark S Granovetter. 1973. The strength of weak ties. *American journal of sociology* 78, 6 (1973), 1360–1380.
- [13] Peter Gräsch, Alexander Felfernig, and Florian Reiffrank. 2013. Recommend: Towards critiquing-based recommendation with speech interaction. In *Proceedings of the 7th ACM Conference on Recommender Systems*, 157–164.
- [14] Frank E Harrell. 2015. Ordinal logistic regression. In *Regression modeling strategies*. Springer, 311–325.
- [15] Jonathan L Herlocker, Joseph A Konstan, and John Riedl. 2000. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*. ACM, 241–250.
- [16] Daniel Herzog and Wolfgang Würndl. 2019. A User Study on Groups Interacting with Tourist Trip Recommender Systems in Public Spaces. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*. ACM, 130–138.
- [17] Anthony Jameson, Stephan Baldes, and Thomas Kleinbauer. 2003. Enhancing mutual awareness in group recommender systems. In *Proceedings of the IJCAI*.
- [18] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2020. A survey on conversational recommender systems. *arXiv preprint arXiv:2004.00646* (2020).
- [19] Oliver P John, Sanjay Srivastava, et al. 1999. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research* 2, 1999 (1999), 102–138.
- [20] Iris A Junglas, Norman A Johnson, and Christiane Spitzmüller. 2008. Personality traits and concern for privacy: an empirical study in the context of location-based services. *European Journal of Information Systems* 17, 4 (2008), 387–402.
- [21] Bart P Knijnenburg, Alfred Kobza, and Hongxia Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71, 12 (2013), 1144–1162.
- [22] Melinda L Korzaan and Katherine T Boswell. 2008. The influence of personality traits and information privacy concerns on behavioral intentions. *Journal of Computer Information Systems* 48, 4 (2008), 15–24.
- [23] Judith Masthoff. 2011. Group recommender systems: Combining individual models. In *Recommender systems handbook*. Springer, 677–702.
- [24] Kevin McCarthy, Maria Salamó, Lorcan Coyle, Lorraine McGinty, Barry Smyth, and Paddy Nixon. 2006. Group recommender systems: a critiquing based approach. In *Proceedings of the 11th international conference on Intelligent user interfaces*, 267–269.
- [25] AKM Mehdy, Michael D Ekstrand, Bart P Knijnenburg, and Hoda Mehrpouyan. 2021. Privacy as a Planned Behavior: Effects of Situational Factors on Privacy Perceptions and Plans. *arXiv preprint arXiv:2104.11847* (2021).
- [26] Soha A Mohamed, Taysir Hassan A Soliman, and Adel A Sewisy. 2016. A context-aware recommender system for personalized places in mobile applications. *Int. J. Adv. Comput. Sci. Appl* 7, 3 (2016), 442–448.
- [27] Shabnam Najafian, Amra Delic, Marko Tkalcic, and Nava Tintarev. 2021. Factors Influencing Privacy Concern for Explanations of Group Recommendation. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*, 14–23.
- [28] Shabnam Najafian, Oana Inel, and Nava Tintarev. 2020. Someone really wanted that song but it was not me! Evaluating Which Information to Disclose in Explanations for Group Recommendations. In *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion*, 85–86.
- [29] A. Napierala, M. 2012. What Is the Bonferroni correction? <http://www.aaos.org/news/aaosnow/apr12/research7.asp>
- [30] Thuy Ngoc Nguyen. 2017. Conversational group recommender systems. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, 331–334.
- [31] Thuy Ngoc Nguyen and Francesco Ricci. 2017. Dynamic elicitation of user preferences in a chat-based group recommender system. In *Proceedings of the Symposium on Applied Computing*, 1685–1692.
- [32] Robert Plutchik. 2001. The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American scientist* 89, 4 (2001), 344–350.
- [33] Rashmi Sinha and Kirsten Swearingen. 2002. The role of transparency in recommender systems. In *CHI'02 extended abstracts on Human factors in computing systems*, 830–831.
- [34] Nava Tintarev and Judith Masthoff. 2007. A survey of explanations in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*. IEEE, 801–810.
- [35] Janice Y Tsai, Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2010. Location-sharing technologies: Privacy risks and controls. *Isjlp* 6 (2010), 119.
- [36] Zhu Wang, Xingshe Zhou, Zhiwen Yu, Haipeng Wang, and Hongbo Ni. 2010. Quantitative evaluation of group user experience in smart spaces. *Cybernetics and Systems: An International Journal* 41, 2 (2010), 105–122.
- [37] Luyan Xu, Xuan Zhou, and Ujwal Gadiraju. 2020. How Does Team Composition Affect Knowledge Gain of Users in Collaborative Web Search?. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, 91–100.

Visual-Meta Appendix

The data below is what we call Visual-Meta. It is an approach to add information about a document to the document itself, on the same level of the content (in style of BibTeX). It is very important to make clear that Visual-Meta is an approach more than a specific format and that it is based on wrappers. Anyone can make a custom wrapper for custom metadata and append it by specifying what it contains: for example @dublin-core or @rdfs.

The way we have encoded this data, and which we recommend you do for your own documents, is as follows:

When listing the names of the authors, they should be in the format 'last name', a comma, followed by 'first name' then 'middle name' whilst delimiting discrete authors with ('and') between author names, like this: Shakespeare, William and Engelbart, Douglas C.

Dates should be ISO 8601 compliant.

Every citable document will have an ID which we call 'vm-id'. It starts with the date and time the document's metadata/Visual-Meta was 'created' (in UTC), then max first 10 characters of document title.

To parse the Visual-Meta, reader software looks for Visual-Meta in the PDF by scanning the document from the end, for the tag @[visual-meta-end]. If this is found, the software then looks for @[visual-meta-start] and uses the data found between these tags. This was written September 2021. More information is available from <https://visual-meta.info> for as long as we can maintain the domain.

@{visual-meta-start}

@{visual-meta-header-start}

@visual-meta{version = {1.1},

generator = {ACM Hypertext 21},

organisation = {Association for Computing Machinery}, }

@{visual-meta-header-end}

@{visual-meta-bibtex-self-citation-start}

@inproceedings{10.1145/3465336.3475104,

author = {Najafian, Shabnam and Draws, Tim and Barile, Francesco and Tkalcic, Marko and Yang, Jie and Tintarev, Nava},

title = {Exploring User Concerns about Disclosing Location and Emotion Information in Group Recommendations},

year = {2021},

isbn = {978-1-4503-8551-0},

publisher = {Association for Computing Machinery},

address = {New York, NY, USA},

url = {<https://doi.org/10.1145/3465336.3475104>},

doi = {10.1145/3465336.3475104},

abstract = { Recent research has shown that explanations serve as an important means to increase transparency in group recommendations while also increasing users' privacy concerns. However, it is currently unclear what personal and contextual factors affect users' privacy concerns about various types of personal information. This paper studies the effect of users' personality traits and preference scenarios --- having a majority or minority preference --- on their privacy concerns regarding location and emotion information. To create natural scenarios of group decision-making where users can control the amount of information disclosed, we develop TouryBot, a chat-bot agent that generates natural language explanations to help group members explain their arguments for suggestions to the group in the tourism domain. We conducted a user study in which we instructed 541 participants to convince the group to either visit or skip a recommended place. Our results show that users generally have a larger concern regarding the disclosure of emotion compared to location information. However, we found no evidence that personality traits or preference scenarios affect privacy concerns in our task. Further analyses revealed that task design (i.e., the pressure on users to convince the group) had an effect on participants' emotion-related privacy concerns. Our study also highlights the utility of providing users with the option of partial disclosure of personal information, which appeared to be popular among the participants.},

numpages = {10},

keywords = {explanation; privacy concern; information privacy; group recommendation},

location = {Virtual Event, USA},

series = {HT '21},

vm-id = {10.1145/3465336.3475104} }

@{visual-meta-bibtex-self-citation-end}

@{visual-meta-end}