

User Modeling for Privacy-preserving Explanations in Group Recommendations

Najafian, S.

DOI

[10.4233/uuid:8197cbd9-64e2-4eae-b424-f2af9ad83641](https://doi.org/10.4233/uuid:8197cbd9-64e2-4eae-b424-f2af9ad83641)

Publication date

2023

Document Version

Final published version

Citation (APA)

Najafian, S. (2023). *User Modeling for Privacy-preserving Explanations in Group Recommendations*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:8197cbd9-64e2-4eae-b424-f2af9ad83641>

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.

User Modeling for Privacy-preserving Explanations in Group Recommendations

User Modeling for Privacy-preserving Explanations in Group Recommendations

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen
chair of the Board for Doctorates
to be defended publicly on
Wednesday 1 February 2023 at 15:00 o'clock

by

Shabnam NAJAFIAN

Master of Science in Informatics, The Technical University of Munich, Germany
born in Zanzan, Iran

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. ir. G.J.P.M. Houben,	Delft University of Technology, promotor
Prof. dr. N. Tintarev,	Delft University of Technology, promotor

Independent members:

Prof. dr. ir. J.F.M. Masthoff,	Utrecht University
Prof. dr. ing. J. Ziegler,	University of Duisburg-Essen, Germany
Prof. dr. ir. A. Bozzon,	Delft University of Technology
Dr. C. Bauer,	Utrecht University
Prof. dr. M.M. Specht,	Delft University of Technology, reserve member



SIKS Dissertation Series No. 2023-03

The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

Published and distributed by: Shabnam Najafian

E-mail: shabnam.najafian@gmail.com

Keywords: Explanation, Group recommendation, Information privacy, Personal information

Printed by: Proefschriftenprinten.nl

Cover by: [Dr. David Maxwell](#)

Typeset in *Manjari* and *Comfortaa* under the terms of the [Open Font License](#)

Style: TU Delft House Style, with modifications by Moritz Beller

<https://github.com/Inventitech/phd-thesis-template>

Copyright © 2023 by S. Najafian

ISBN: 978-94-6366-640-4

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

An electronic version of this dissertation is available at

<http://repository.tudelft.nl/>.

To all those free-spirited human beings who have fought for justice. To my people in Iran and their longing for liberty. To Iranian women and their uncommon bravery in their fight against oppression, whom I'm proud to share a fraction of my identity with them.

Az yaşa, Azad yaşa, İnsan yaşa. (Life's short, live free, remain humane.)

Babək Xorrəmdin (Babak Khorramdin)

Contents

Acknowledgments	xi
1 Introduction	1
1.1 Research Motivation	2
1.2 Research Questions	3
1.3 List of Publications	5
2 Literature Review	7
2.1 Group Decision Making	8
2.1.1 Aggregation Strategies	8
2.1.2 Group Deliberation	9
2.2 Explanations for Groups	10
2.3 Privacy in Group Explanations	11
2.3.1 Individual Characteristics	12
2.3.2 Situational Characteristics	12
2.4 The Trade-off between Risk and Benefit on Information Disclosure	13
2.4.1 Antecedents of Information Disclosure	14
2.4.2 Antecedents of Disclosure Benefits	14
2.4.3 Antecedents of Privacy Risk	15
2.5 Conclusions	17
3 Deciding on the Content for Group Explanations	19
3.1 Experiment 1: Aggregation Strategies as a Basis for Group Explanation	21
3.1.1 Preliminary Definitions: Aggregation Strategies	21
3.1.2 Experimental Design	24
3.1.3 Results	28
3.1.4 Discussion	33
3.2 Experiment 2: Aggregation Strategies versus Explanations	35
3.2.1 Experimental Design	36
3.2.2 Results	41
3.2.3 Discussion	43
3.3 Experiment 3: Formulating Group Explanations	47
3.3.1 Preliminary Definitions: Explanation Styles	47
3.3.2 Experimental Design	48
3.3.3 Results	51
3.3.4 Discussion	52
3.4 Chapter Conclusions	55

4	Privacy in Group Explanations	59
4.1	Experiment 1: Privacy Preferences	61
4.1.1	Preliminary Definitions: Generating Natural Language Explanations	62
4.1.2	Experimental Design	63
4.1.3	Results	66
4.1.4	Discussion	73
4.2	Experiment 2: Factors that Affect Privacy	76
4.2.1	Experimental Design	76
4.2.2	Results	81
4.2.3	Discussion	85
4.3	Experiment 3: Disclosing Location and Emotion Information.	88
4.3.1	Experimental Design	89
4.3.2	Results	95
4.3.3	Discussion	99
4.4	Chapter Conclusions	102
5	Antecedents of Information Disclosure	105
5.1	Experiment 1: Influence of Task Design on Information Disclosure.	107
5.1.1	Experimental Design	108
5.1.2	Results	108
5.1.3	Discussion	109
5.2	Experiment 2: The Trade-off between Risk and Benefit on Information Disclosure	111
5.2.1	Experimental Design	112
5.2.2	Results	120
5.2.3	Discussion	124
5.3	Chapter Conclusions	127
6	Conclusions and Future Work	131
6.1	Research Questions Revisited	132
6.2	Practical Considerations	134
6.2.1	Empirical Implications	135
6.2.2	Methodological Lessons Learned	136
6.3	Limitations and Future Directions	136
6.4	Summary.	138
	Bibliography	141
	Summary	157
	Samenvatting	159
	Curriculum Vitæ	163
	SIKS Dissertation Series	165

Acknowledgments

This Ph.D. thesis is the culmination of the effort and support of several people to whom I am incredibly grateful. First of all, I would like to express my highest gratitude to my promoters, Geert-Jan Houben and Nava Tintarev. Geert-Jan, thank you for your constructive and critical comments during the insightful discussion. You were constantly pushing all of us in the WIS group to think about the broader picture. Thank you, Nava, for your scientific guidance, constructive criticism, and thoughtful remarks. The nudges you consistently gave me to push my research and writing one more step were essential to the completion of this dissertation.

I would like to thank the members of my defense committee: Prof. dr. ir. J.F.M. Masthoff, Prof. dr. ing. J. Ziegler, Prof. dr. ir. A. Bozzon, Dr. C. Bauer, and Prof. dr. M.M. Specht, for accepting to be part of my committee and for their time, patience with me, and eagerness to help.

I had the privilege to work with a number of remarkable researchers: Marko Tkalcic, Jie Yang, Amra Delic, Daniel Herzog, and Bart Knijnenburg. Thank you for the valuable brainstorming sessions, the fruitful collaboration, and the excellent results. Special thanks to Dr. Knijnenburg, who played a significant role in guiding me and answering my seemingly never-ending questions, particularly regarding statistics, through my last paper. My thanks also go to my master's thesis supervisor, Dr. Wolfgang Wörndl, whose mentorship and thoughtful remarks greatly influenced my interest in research and the scientific method.

Many thanks to the current and former members of the Web information systems group at the Delft University of Technology: Marcus Specht, Alessandro Bozzon, Christoph Lofi, Claudia Hauff, Jie Yang, Ujwal Gadiraju, Avishek Anand, Maria Soledad Pera, Rihan Hai, Yue Zhao, Achilleas Psyllidis, Dimitrios Bountouridis Carlo van der Valk, Pavel Kucherbaev, Sara Salimzadeh, Georgios Siachamis, Peide Zhu, Guanliang Chen, Andra Lonescu, Mónica Marrero, Arthur Camara, Vincent Gong, Gustavo Penha, Marios Fragkoulis, Ziyu Li, Nirmal Roy, Pedro Fortunato Silvestre, Mohammad Khalil, Lixia Chu, Lorenzo Corti, Garrett Allen, Tahir Abbas, Gaole He, Lijun Lyu, Tamara Brusik, and Daphne Stephan. I am particularly grateful to Oana Inel, Mesut Kaya, Tim Draws, Alisa Rieger, Francesco Barile, and Amir Ebrahimi Fard – my colleagues currently or formerly working in the Epsilon team for their constructive feedback. I would like to thank my wonderful office mates: Sepideh Mesbah, Shahin Sharifi, Sihang Qiu, Ioannis Petros Samiotis, Christos Koutras, Agathe Balayn, Andrea Mauri, and Manuel Valle Torre, for being people that are a joy to be around and for happy moments we shared. To Asterios Katsifodimos for always being energetic and positive, accepting to be my second supervisor for a time, and teaching me how to cope in challenging circumstances. To Sepideh Mesbah for all the enjoyable talks and breaks we had. I also want to thank her for motivating and supporting me at moments when I wanted to give up and for listening all my complaints. To Felipe

Moraes for offering good advice when I needed it. To David Maxwell for his meticulous and delicate artwork on my thesis cover.

I really appreciate spending time with my dear friends. Special thanks must go to Marzieh, Mehdi, Saba, Sina, Maryam, Ehsan, Sepideh, Shahin, Taraneh, Delaram, Nima, Fatima, Erfan, Mahsa, Saeed, Sepideh, Farshid, Shokoufe, Milad, Zahra, Mostafa, Mojde, Soheib. And their little ones, Avid, Yara, and Mandana. Those memorable weekends we spent together were a great boost to my energy. Shima, I am so happy to have you as a part of my life; thank you for always being so kind and caring to me. Thoughts of each other will not be hindered by distance. We always feel welcomed at your house when we are back in Iran by both you and Shafagh. Laleh, thank you for challenging me to be my best self. I greatly admire you as a very strong woman and a role model. Farshad, thank you for being patient and supporting me when I had coding questions. Samira, back in Iran, we had so much fun together. Additionally, receiving new year cards every year from you and Emad is so heartwarming.

It would be impossible for me to survive without the support of my entire family. You are all so amazing and do everything in your power to support me, regardless of where or what I am doing. My parents, Baba and Pari joon, for their unconditional love and for always being my biggest cheerleaders. My brother Ali and sister-in-law, Nafiseh, thank you for your care and encouragement. Ali, I appreciate you buying me books from my childhood that inspired me; I learned a lot from you. To Maryam and Alireza, my sister and brother-in-law, for always being there for me. Maryam, having a sister like you is the best thing anyone could ever ask for. To my adorable niece and nephew, Barin and Noyan, the apples of my eye and the cheer of my life, I am so lucky to have you. Last but certainly not least, I wish to express my deepest gratitude to my better half and my best friend, Majid. I am so grateful for the sacrifices you made for me during my Ph.D., especially by traveling long distances from Germany to the Netherlands, back and forth. The past four years have been full of highs and lows. You have been there for me every step of the way. Thank you for everything, and cheers to many more years of adventures.

I hope I have been able to convey to everyone I mentioned here (and maybe I forgot some names) how much I value you in my life. It is such a blessing to have you all in my life. I appreciate everything you have done for me.

*Shabnam Najafian
Delft, January 2023*

1

Introduction

My thesis investigates what makes good explanations for group recommendations, considering the privacy concerns of group members. Let's give an example. Have you ever been to lunch with other colleagues on a business trip? Do you recall how long it took you to pick a restaurant? In these situations, recommender systems could help people decide, e.g., where to go. Recommender systems are decision support systems helping users to identify one or more items that satisfy their requirements. Most often, recommender systems propose items to individual users. However, there are many scenarios where a group of users will consume a recommendation and need support for group decision-making. A group recommender system is a system that recommends items to groups of users collectively, given their preferences. An example is a system for suggesting places to visit to a group of colleagues traveling together. For example, think of a group decision regarding the next places to visit in a colleagues'/friends' group traveling. Explanations, for such recommendations, in this context, act as complementary information, describing how specific recommendations are generated to help the group make informed decisions on whether to follow or not follow recommendations. However, there are many types of information to include and many ways to formulate an explanation, and it is not clear which information should be shown in the explanation for a group. Besides, explanations for groups are different from explanations for single users in that they should consider the privacy aspect (e.g., people might be sensitive to disclosing some of their information in the group). In this chapter, I first introduce the motivation of this Ph.D. thesis of developing explanations for group recommendations/decisions context. To the best of my knowledge, this thesis is the first work that studied group explanations from the perspective when the privacy aspect is included. Then I list the research questions that guide my thesis to design explanations for groups and summarize the corresponding contributions. This includes studying what information to disclose and what not to disclose in a group explanation and what factors and how influence the decision of information disclosure in a group explanation, e.g., the group members' personality, the relationship between them, whether their opinion is aligned with the majority in the group or not. Finally, I present a list of publications carried out during this thesis.

1.1 Research Motivation

Have you ever been to lunch with other colleagues on a business trip? Do you recall how long it took you to pick a restaurant? Imagine you start walking to one restaurant only to discover that Person A wants to eat Halal, Person B has an auto-immune protocol diet, and Person C prefers a low-budget place. After visiting a restaurant, your group might also need to pick where to go next (e.g., a war museum, a cannabis store, etc.). Will you speak out if your preferences do not align with the majority in the group? What about when you do not have close relationships with other group members? Not only is it challenging to cater to multiple preferences, it can also be difficult to surface individual preferences in order to make an informed group decision!

In many domains, such as tourism [10, 85], people often consume recommendations in groups rather than individually. There are different approaches in the literature for considering different preferences within a group. A common approach is aggregation strategies that combine the individual preferences of all group members and predict an item that is suitable for the group [74, 76, 85]. An advantage is that this approach can be easily explained to the group, which can be a good start for designing group explanations. A disadvantage is that they are simple and do not, for example, evolve based on users' needs, like machine learning approaches. There is no best aggregation strategy that exists, and for each recommendation, some individuals might not be happy with the recommendation [2]. For example, the Fairness Strategy (an aggregation strategy) [74] might recommend an item that one or more group members do not like but will recommend at other times other items that they do like, in order to compensate.

In these situations, explanations can clarify such trade-offs, help people comprehend how these recommendations are generated, make it easier to accept items they do not like, and ultimately facilitate reaching a group consensus and making informed decisions [7, 30, 84, 131]. Explanations can be regarded as additional information (i.e., in this thesis, this information is textual only) that accompanies the recommendations and serves various goals, such as increasing satisfaction (the ease of usability or enjoyment of the used recommender system) [124]. Many studies have demonstrated the benefits of adding explanations to automated recommendations (e.g., [42, 117]). Previous research in this area has focused on explaining individual recommendations [42, 117]. Similar to explanations for single-user recommendations, explanations for groups can be designed based on the underlying recommendation algorithm. However, in the context of group recommendations, formulating explanations is even more challenging as other aspects must be considered. To study what to explain in a group explanation, we initially considered explaining the aggregation strategies as central. Then we discovered that other factors such as privacy and group composition (e.g., whether the user's preferences align with the majority in the group, their social relationship in the group) were much more influential. Explaining why certain items are recommended can help users agree on a joint decision within a group [30, 98], but the value of such explanations should be traded-off against the desire to preserve individuals' privacy by not disclosing information they do not want to disclose in that group. For example, if other group members know each other's preferences, it is easier to reach a consensus/converge on a decision and find something they all want to do. Still, there might be personal information/preferences one does not wish to disclose. So the explanation should strike a balance between giving the group enough information to

achieve what they want without revealing the information they are uncomfortable with.

The few existing works on generating explanations for group recommendations primarily consider the need for transparency [30, 108]. However, considering privacy aspects of explanations for groups of users is essentially an open research issue. To the best of my knowledge, this thesis is the first work that studied group explanations from the perspective when the privacy aspect is included. My ultimate purpose is to help people with better group decision-making by providing them with privacy-preserving explanations to help group members explain their arguments for or against the suggested items (places) to the group. To this end, I look at both static (Chapter 3 and part of Chapter 4) and interactive information provision (Chapter 5 and part of Chapter 4) for the group explanations. This thesis is mainly conducted in the context of *tourism*—a domain that is suitable for studying group decisions, as it is relatable for many participants and commonly involves coordinating with a group of people. In order to accomplish this goal, we need a more fundamental understanding of factors that influence people’s disclosure decisions in various group decision-making contexts. I, therefore, conducted a number of studies to contextualize users’ disclosure decisions to understand which individual and situational factors need to be considered in order to predict whether users are willing to disclose certain personal information to help with group decision-making or not. Research in the online privacy context shows most Internet users trade-off the anticipated benefits with the risks of disclosure to decide on their information disclosure. In making this trade-off, these users decide to disclose how much, if any, information is requested from them. Based on this deeper understanding of users’ disclosure behavior, the core contribution of this thesis is a privacy disclosure model containing different individual and situational models/characteristics that help to predict users’ disclosure intention in a group decision/recommendation context.

1.2 Research Questions

This thesis investigates the following main research question in the group recommendation/decision context.

How do different factors influence individual group members’ requirements towards a group explanation?

To answer our main research question, we organized the work into three research sub-questions (the research questions have evolved in the process).

RQ1 What information should be disclosed in a group explanation to increase group members’ satisfaction?

As a first step, in Chapter 3, we started evaluating with people, which kinds of explanations are meaningful to include in a group recommendation to increase group members’ satisfaction (**RQ1**). We contributed novel suggestions for formulating explanations for group recommendations (the explanations are textual and static). Inspired by explanations for single-user recommendations, explanations for groups included information about the underlying recommendation algorithm (in this case, aggregation strategies). However, we

decided not to proceed with aggregation strategies as they do not show a significant effect in, e.g., group satisfaction or consensus. The user comments suggested that other factors, such as privacy (people might be sensitive to disclosing some of their information), were much more influential in the studied settings. This leads us to **RQ2** as follows. As we will see, privacy risk will be a recurring theme in the remainder of the thesis. We also highlighted the challenges of benchmarking and replicating studies in the context of group recommendations and explanations.

RQ2 How do different factors (i.e., individual differences, group dynamics, etc.) influence individual group members' privacy risk perception of information disclosure in a group explanation?

As mentioned above, the user comments highlight the need to consider privacy in designing a group explanation. This leads us to the second research question. In Chapter 4, we investigate the factors that one should model in the group to consider group members' privacy risks regarding information disclosure in a group explanation (**RQ2**). We investigated some factors identified in the literature that influence individual privacy risks of information disclosure in our case in a group explanation, i.e., group members' personality, the type of relationship they have in the group, and preference scenario (whether their preferences are aligned or not aligned with the preferences of the majority in the group). We gave design recommendations for automatically generating group explanations in this context (e.g., when the recommended item does not align with the majority preferences, one should be cautious about disclosing the identity of the people with the minority preferences in the group, together with their strong opinions). We saw that it is not enough to only look at privacy risks for predicting people's disclosure behavior. Inspired by literature, next, in **RQ3**, we study how people trade-off the anticipated disclosure benefit with the privacy risk of disclosing information to decide how much, if any, information to disclose to be able to model their privacy disclosure.

RQ3 How do people trade-off between disclosure benefit versus privacy risk of information disclosure in a group explanation?

So, in Chapter 5, we focused on **RQ3**, and we studied how people trade-off the anticipated disclosure benefit with the privacy risk of disclosing information to decide how much, if any, information to disclose in a group explanation. In such decisions in group recommendations, users face a dilemma: they want to enjoy the benefits that may result from sharing or disclosing information with other group members (e.g., support their arguments about what places to visit or to avoid), but they also want to reduce the risk that this data may have (e.g., leaving a negative impression on others). In making this trade-off, users usually decide to disclose some but not all information that is requested from them [62]. The findings of our study regarding user disclosure decisions can be utilized to automatically predict a proper fit between users' desire for privacy and their need for transparency to make better group decisions.

Study Platform. In Chapters 4 and 5, to answer the research questions, I developed an open-source web-based chat-bot, to study group disclosure decisions and evaluate explanations for group recommendations. To create realistic scenarios of group decision-making

where users can control the amount of information disclosed in the group and have iterative interaction between group members, I developed TouryBot. This chat-bot agent generates natural language explanations to help group members explain their suggestions to the group. This publicly available implementation can be easily used and adapted for empirical studies in a group recommendation/decision context. The code is available at the following address <https://osf.io/z3hnp/>.

Note all our experiments received ethical committee approval from the Human Research Ethics Committee (HREC) at TU Delft.

1.3 List of Publications

A complete list of publications on which the research chapters were based is presented below:

Chapter 3 is based on a conference paper, a workshop paper, and a late-breaking results paper.

- Shabnam Najafian, Daniel Herzog, Sihang Qiu, Oana Inel, and Nava Tintarev. You do not decide for me! Evaluating explainable group aggregation strategies for tourism. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, pages 187–196, 2020.
- Francesco Barile, Shabnam Najafian, Tim Draws, Oana Inel, Alisa Rieger, Rishav Hada, and Nava Tintarev. Toward benchmarking group explanations: Evaluating the effect of aggregation strategies versus explanation. 2021.
- Shabnam Najafian and Nava Tintarev. Generating consensus explanations for group recommendations. In *UMAP Latebreaking results*, 2018.

Chapter 4 is based on a workshop paper and two conference papers.

- Shabnam Najafian, Oana Inel, and Nava Tintarev. 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*, pages 85–86, 2020.
- Shabnam Najafian, Amra Delic, Marko Tkalcic, and Nava Tintarev. Factors influencing privacy concern for explanations of group recommendation. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*, pages 14–23, 2021.
- Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. Exploring user concerns about disclosing location and emotion information in group recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, pages 155–164, 2021.

Chapter 5 is based on the part of a conference paper and a journal paper (under review).

- Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalčić, Jie Yang, and Nava Tintarev. Exploring user concerns about disclosing location and emotion information in group recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, pages 155–164, 2021
- Shabnam Najafian, Geoff Musick, Bart Knijnenburg, and Nava Tintarev. How do People Make Decisions in Disclosing Personal Information in Tourism Group Recommendations in Competitive versus Cooperative Conditions? *Modeling and User-Adapted Interaction (UMUAI) Journal*, 2022 (under review).

The thesis also benefited from insights gained from the following workshop paper:

- Öykü Kapcak, Simone Spagnoli, Vincent Robbmond, Soumitri Vadali, Shabnam Najafian, and Nava Tintarev. Tourexplain: A crowdsourcing pipeline for generating explanations for groups of tourists. In *Workshop on Recommenders in Tourism co-located with the 12th ACM Conference on Recommender Systems (RecSys 2018)*, volume 2222. CEUR, 2018

2

Literature Review

This thesis is about studying what makes good explanations for group recommendations. In this chapter, I illustrate the relevant literature in this regard. Making a joint decision in the group, i.e., deciding where to visit with the group of colleagues on a business trip, is challenging as people often have different preferences. Various strategies can be used to predict an item suitable for the group collectively, given the preferences of all members. Each such aggregation strategy, however, has its trade-offs. In this situation, explanations can help people understand why certain items are recommended i.e., to increase group members' satisfaction. However, among many ways of formulating an explanation for groups, it is unclear which information needs to be shown in the explanation. Inspired by explanations for single-user recommendations, explanations for groups included information about the underlying recommendation algorithm (in this case, aggregation strategies). This chapter introduces the aggregation strategies, which was initially included in our designed group explanations. Besides, I discuss conversational interfaces as an interactive group recommendation approach, which I use to study and model group members disclosure behavior in the group recommendation context. Then I provide an overview of existing research related to explanations in group recommender systems and discuss privacy aspects in explanations for groups, which arise in these scenarios. I finish this chapter by summarizing the effects of disclosure benefit versus privacy risk and their antecedents (i.e., group member's personality) on information disclosure in a group explanation.

2.1 Group Decision Making

Group recommender systems (GRSs) have been developed to support group decision-making processes with recommendations expected to satisfy a group of people, not just a single person. Most of the previous research has assumed that only by aggregating individual preferences can the system come up with a satisfying recommendation for all individuals in the group. Following, I introduce the existing aggregation strategies as a static group recommendation approach which I initially include their descriptions in our proposed group explanations. Besides, I discuss conversational interfaces as an interactive group recommendation approach, which I use to study and model group members disclosure behavior in the group recommendation/discussion context. In an interactive group recommender system people can a) disclose their personal information to explain and support their arguments about what recommended items to accept or to avoid (e.g., this place is too expensive for my budget) and b) protect their privacy by not disclosing too much.

2.1.1 Aggregation Strategies

There are two main approaches to generate group recommendations: (i) aggregated predictions or strategies, that aggregate individual recommendations (item-ratings predictions) and recommend items with the highest aggregated scores to the group; or (ii) aggregated models, which instead of aggregating recommendations (item-ratings predictions) for individual users, this approach construct a group preference model (group profile) that is then used for determining recommendations [30]. Several aggregation strategies inspired by Social Choice Theory have been proposed to aggregate individuals' information [76]. An overview of these strategies, known as social choice-based aggregation strategies, can be found in Masthoff [74]. Following, I describe in detail six of the most utilized social choice-based aggregation strategies: two consensus-based aggregation strategies, *Additive Utilitarian (ADD)* and *Fairness (FAI)*, one majority-based strategy, *Approval Voting (APP)*, and three borderline strategies, *Least Misery (LMS)*, *Majority (MAJ)*, and *Most Pleasure (MPL)*, according to the categorization in [76, 114].

- *Additive Utilitarian (ADD)*: is a consensus-based strategy [114], so it takes into account the preferences of all group members. ADD recommends the item with the highest sum of all group members' ratings. Applying this strategy to the example given in Table 2.1, restaurant B (Rest B) will be recommended to the group as the sum of the ratings of all members for restaurant B is 13 which is higher than other items.
- *Fairness (FAI)*: is a consensus-based strategy [76] well suited in the context of repeated decisions. In FAI the items are ranked as the individuals are choosing them in turn. Applying this strategy to the example given in Table 2.1, by assuming that Anna is the next to choose, restaurant B will be recommended to the group, since it achieves the highest rating for her.
- *Approval Voting (APP)*: is a majority-based strategy [114], so it focuses on the most popular items among group members. APP recommends the item which has the highest number of ratings that are greater than a predefined threshold, e.g., 3. Applying this strategy to the example given in Table 2.1, restaurant B will be recom-

mended to the group three out of four group members Anna, Sam, Leo gave it ratings higher than 3.

- *Least Misery (LMS)*: is a borderline strategy [114], so it takes into account only a subset of group members' preferences. LMS recommends the item which has the highest of all lowest ratings. Applying this strategy to the example given in Table 2.1, restaurant A will be recommended to the group as Alex and Anna gave it a rating of 2 which is the highest rating among lowest ratings regarding items (the other ratings are 1s).
- *Majority (MAJ)*: is a borderline strategy [114] which recommends the item with the highest number of all ratings representing the majority of item-specific ratings. Applying this strategy to the example given in Table 2.1, restaurant B will be recommended to the group as 3 out of 4 group members gave it a high rating.
- *Most Pleasure (MPL)*: is a borderline strategy [114] which recommends the item with the highest of all individual group members ratings. Applying this strategy to the example given in Table 2.1, restaurant C will be recommended to the group as Alex gave it the rating of 5, which is the highest ratings among all items' high ratings.

Table 2.1: Ratings of group members for the restaurants (1: the worst, 5: the best) from Tran et al. [131].

	Alex	Anna	Sam	Leo
<i>Rest A</i>	2	2	4	4
<i>Rest B</i>	1	4	4	4
<i>Rest C</i>	5	1	1	1

Social choice-based aggregation strategies are widely used in the group recommenders literature [76]. An advantage is that this approach can be easily explained to the group, which can be a good start for designing group explanations. A disadvantage is that they are simple and do not, for example, evolve based on users' needs, like machine learning approaches. In Masthoff [76], several experiments are presented to identify the best strategy in terms of perceived group satisfaction. The results, however, show that there is no winning strategy (i.e., satisfy all group members), but different strategies perform well in different scenarios (i.e., based on the level of difference in group members preferences).

2.1.2 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 [48]. 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 [91]. Therefore, previous research has introduced strategies to enable interaction between group members during the process [47, 79].

Conversational interfaces specifically have been shown to lead to higher satisfaction of the user, requiring less interaction and increasing the likelihood of them using the system in the future [39]. In this direction, Nguyen and Ricci [92] presented 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 recommendations to each member, in our system, the chat-bot suggests what is best for the group based on an aggregation strategy 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.

2.2 Explanations for Groups

Although there exist many studies on group recommendations, only a few of them focus on generating explanations in group recommendation contexts. Different from single-user recommender systems, the role of explanations for groups is even more challenging, as multiple functions need to be met, besides explaining why certain items are recommended [30, 98] – to help users agree on a common decision, improve users’ perceived fairness, perceived consensus, and satisfaction [30, 131], as well as preserve their privacy by not disclosing information they are not comfortable with (see Chapter 4).

The generation of explanations for group recommendations depends on how the group recommendations were generated in the first place. One approach to generate group recommendations, called aggregated predictions. It aggregates individual item-ratings predictions and recommends to the group items with the highest aggregated scores [30]. Explanations based on this approach reveal the underlying mechanisms of the employed social choice-based preference aggregation strategies [131].

Apart from their styles, explanations can be represented in different ways, e.g., as textual representations, or as graphical representations. The most frequent way of presenting explanations is by far Natural language generation (NLG). NLG is a sub-field of artificial intelligence and computational linguistics used for producing understandable texts in English or other human languages from a given set of text or data [109]. This also includes explanations that are based on pre-defined templates which were, for example, instantiated with lists of features before they were presented to the user [99]. In this thesis, I also use a template-based NLG technique by adding pre-defined templates which can be easily adapted and extended based on each individual group member selected options. Following are some examples of recent research on recommender systems that have focused on generating personalized natural language explanations for group recommendations.

Quijano-Sanchez et al.[108] have focused on using group’s social factors (e.g. users’ personal relationships within a group) for explanation generation and proposed Personalized Social Individual Explanation for groups of users. Their goal was to justify the recommended item for the group’s welfare. They relied on users’ sense of justice and social bonds to help them comprehend why the recommender has presented a specific item as the best option for the group. , i.e., avoid explanations that might damage friendships (using tactful explanation). *“Although we have detected that your preference for this item is not very*

high, your close friend X (who you highly trust) thinks it is a very good choice.” They showed adding a social component to explanations enhances the impact that explanations have on users’ likelihood to follow the recommendations and consequently increases the system’s persuasiveness, efficiency, trustworthiness, and usability. And that also, the more social factors that are included in the explanation the better perception of the received group recommendation [108].

Recently, Tran et al. [131] proposed three types of textual explanations taking into account group dynamics aspects, such as fairness, consensus, and satisfaction of users with group recommendations. Type 1 was based on preference aggregation strategies, Type 2 is decision history in addition to Type 1, and Type 3 is future decision plans in addition to Type 1. For instance, a Type 2 explanation for Additive Utilitarian strategy is as follows: “Item X has been recommended to the group since it achieves the highest total rating. This decision supports the preferences of users u_a , u_b , and u_c who were treated less favorably in the last n decisions” [131].

The existing works on generating explanations for group recommendations primarily consider the need for transparency, e.g., to clarify the reasoning and data behind a recommendation to help users better understand how the recommender system works and why a specific item has been recommended [128]. However, when generating explanations for groups, privacy becomes of great relevance as well. The work of Herzog and Wörndl [43] highlights the need for privacy in a group context, as it found that there was a greater amount of interaction with distributed displays when people were interacting on their individual devices rather than on a shared public display. For example, when using a shared public display to enter sensitive data, privacy was a concern for many participants. We will discuss privacy in greater detail in the next section.

2.3 Privacy in Group Explanations

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 [124]. However, when generating explanations for groups rather than individuals, privacy becomes of great relevance.¹ I investigated which information people would like to disclose in explanations for group recommendations (see Section 4.1). I extended the work to evaluate the factors that have an impact on privacy concerns for group recommendations (see Section 4.2). We show an impact deriving from (i) the *personality*, (ii) the *preference scenario* and (iii) the *relationship type*. Furthermore, Mehdy et al. [80] suggested to consider the *information type* when modeling users’ situation-specific privacy concerns. In the following subsections, we discuss relevant literature on both individual characteristics (e.g., group member’s personality) and situational characteristics (e.g., preference scenario, type of the relationship, information type, and task design).

¹According to my knowledge, this thesis is the first to analyze group explanations in the context of privacy issues. Hence, I included some self-reference in this section to motivate the opted path for the literature review.

2.3.1 Individual Characteristics

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 [14]. Bansal et al. [5] 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. [52] showed significant effects of Agreeableness but suggesting a negative effect (i.e., more agreeable people were less concerned about their privacy). Related work shows mainly that the three personality traits *agreeableness*, *extraversion*, and *neuroticism* are related to privacy concern [5, 52, 65].

2.3.2 Situational Characteristics

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 [3, 32]. This was confirmed in Section 4.2, which showed that people having minority preferences expressed higher privacy concerns, in particular for the information related to their emotions.

Type of the Relationship

Social relationships have been shown to be a contextual factor that has an impact on privacy concerns in information sharing [25, 38, 43, 80]. 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) [80]. Additionally, Wang et al. [135] 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, we showed that privacy concerns are perceived more in loosely-coupled heterogeneous groups than tightly-coupled homogeneous (see Section 4.2). For the remaining of the thesis, I therefore focus on loosely coupled (weak ties) heterogeneous groups to consider privacy concern in an extreme case.

Information type

By information type here I refer to user's personal information type. Existing studies have suggested that the magnitude of privacy concerns depends on the type of information to disclose. They found that privacy concern varies depending on information types [80]. In this dissertation, I 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 [82], while our results highlighted location and emotion information as the information that generates the higher privacy concerns in the context of group recommendations in tourism (see Section 4.2).

Previous work also highlights that it is not just the type of information that is sensitive to disclose, but the granularity at which it is disclosed [132]. For example, 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. [13] 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.

Caliskan Islam et al. [9] categorized emotion as private information. Graham et al. [37] 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. To decide which personal information to include in the study, we used private information categories listed in Caliskan et al. [9], which derived from users' tweets on Twitter, and personal information used in Knijnenburg et al. [62], in an online health application context. We included those that are relevant to tourism recommended systems context, namely the following personal information: *emotion-related information*, *location-related information*, *financial-related information*, *religion-related information*, *health-related information*, *sexuality-related information*, and *alcohol-related information*.

Task design

In the literature there are studies focuses on competitive vs cooperative incentives for sharing information. Importantly, this information sharing is NOT personal information sharing. But inspired by this body of research for general information sharing we look at personal information disclosure specifically. Toma and Butera [129], stated that competition activates the fear of being exploited (risk vulnerability), but also the desire to exploit other people. They also add, in all information exchange situations, competition activates tactical deception tendencies aimed at maintaining a positive self in other people's eyes [129]. When they are doing it for the group (e.g., in the cooperative task) then it seems their own privacy risk becomes less important.

Privacy concerns itself consists of a calculus of privacy risk and disclosure benefit which I cover in the next section.

2.4 The Trade-off between Risk and Benefit on Information Disclosure

When people are in the situation where they have to decide where to visit next in a group traveling, they have to make a trade-off between disclosing their personal information to

explain and support their arguments about where to visit or where to skip (i.e., this place is too expensive for my budget), while not violating their privacy by disclosing too much, especially given who those group members are.

So in this section, we discuss relevant literature on what affects this trade-off to disclose personal information in group recommendation context.

2

2.4.1 Antecedents of Information Disclosure

As one of the most prominent information privacy research frameworks, the *privacy calculus theory* examines information disclosure as a decision in which people trade off risks against benefits [18]. In the privacy calculus framework, perceived privacy risk is the degree to which people believe there is a potential for loss associated with the release of personal information [22] and benefits are the context-specific gains individuals expect in exchange for the information they provide [51].

Disclosure benefit. People may respond differently to information disclosure based on their assessment of inherent trade-offs of risks and benefits. In our context (tourism group decisions/recommendations), perceived disclosure benefits refer to the extent to which users believe disclosing their personal information to their group members is beneficial for the group decision or for their own negotiation position within the group. If the users feel that they get some benefits, then they will give up some level of their privacy in return for the perceived benefits [12, 56, 118, 136].

Privacy risk. On the other hand, perceived risks include all the problems and difficulties that the users might face when the other parties have access to their personal information. Perceived privacy risk in our context can be defined as the “expectation of losses associated with the disclosure of personal information in the group”, adapted from Xu et al. [137]’s definition for online providers. Therefore, if users perceive that they are at risk when they disclose their personal information, this can decrease their willingness to share information with online providers [55, 73, 94]. For example, Keith et al. [55] found that increased perceived privacy risk from a mobile application decreases users’ intention to share personal information, including location and financial information.

As can be seen, there is a tension between the perceived benefit of disclosing personal information and the degree of risk individuals perceive by disclosing their information in the group: depending on the situation, if people find that the benefit of disclosing their information outweighs the involved risk, they will disclose the information. Otherwise, they will not disclose their information in the group.

2.4.2 Antecedents of Disclosure Benefits

Perceptions of benefit can be affected by different factors. Milne et al. [81], delineate cost-benefit perceptions of information exchange and indicate that some consumers do not mind revealing private information to a company if they receive specific benefits for providing the information. The benefit is context-dependent, and one’s evaluation of benefit is influenced by a) the amount of *trust* the individual has in the receiver (or, in our group recommender context, the group) [111, 115], b) the *preference scenario* (having minority or majority preferences compared to other group members, see Section 4.2), and c) the

task design (whether group members were instructed to convince other group members of their opinion, or not) [129, 130]. For example, if an individual is in the minority position, disclosing more information may help support their arguments to the group compared to an individual in the majority who does not have to take that effort. I address each of these in turn:

Trust. Trust has mainly been studied in individual contexts, e.g., trust in an app or an institution to which users disclose their information. In such contexts, Kehr et al. [54] showed that trust positively affects the perceived benefits of disclosing information. We assume this can be similar to trust in a group with whom one travels. For example, when group members trust the other individuals in the group then they will perceive a lower risk, and hence greater benefits in providing their personal information [111, 115].

Preference scenario. The “preference scenario” represents whether the active user’s preferences are in the minority or majority within the group. People whose preferences are in the minority may perceive more benefit from providing the group with reasons behind their preferences than those whose preferences are in the majority. In this thesis, I consider triads (a group containing three members) to explore this parameter. I mention “both majority scenarios”, because when the social positions of the group members are not equal, the majority scenario itself can also have two conditions, depending which other member has the same preference as the participant. In this thesis, I consider one other group member to be a peer of the user and the other group member a superior. When the participants’ preferences are in line with their superior and opposite to their peer (which I call “*boss majority*”) this will have different social implications than when the participants’ preferences are in line with their peer and opposite to their superior (which I call “*peer majority*”).

Task design. The competitive or collaborative nature of task design often influences group member behavior and has previously been explored in group decision-making literature [129, 130]. Notably, the competitive mindset often urges group members to share information with the goal of ‘winning’ the discussion to be ‘right’ [45, 130]. This competitive mindset might influence group members to share more information to reach a group decision that matches their preferences.

2.4.3 Antecedents of Privacy Risk

Above, we looked at factors contributing to perceived disclosure benefits—a perception that should increase disclosure. Now we look at factors that contribute to perceived privacy risk—a perception that, in contrast, should decrease disclosure. Risk has been defined as uncertainty resulting from the potential for a negative outcome [41], and one’s evaluation of risk is influenced by a) the amount of *trust* the individual has in the receiver (or, in our group recommender context, the group) [111, 115], b) the *preference scenario* (see Section 4.2), and c) the *task design*. Further, trust is influenced by one’s general *privacy concern* perception [54], and finally one’s general *privacy concern* is influenced by one’s *personality* [5, 52, 65]. I address each of these in turn:

Trust. Trust has been addressed by a number of prior studies and is generally viewed as a type of belief that users can confide on certain entities to protect their personal information [73]. Trust is an important factor that can negate the effects of perceived risk [46, 66]. If trust is established in the mind of the users, then they will perceive a lower risk in providing their personal information [111, 115]. In the context of group decisions/recommendations, when group members trust the other individuals in the group, they are more willing to accept personal vulnerability, which consequently perceives less privacy risk [69, 78]. Previous studies have demonstrated that perceived trust is positively related to reducing the privacy risks of personal information disclosure [68, 90, 116].

Preference scenario. Several studies suggest that the relative preferences of group members (i.e., the preference scenario), could impact the privacy risk. In particular, people whose preferences are in the minority within the group could decide not to share their preferences to match the opinions of the majority, for a phenomenon known as conformity [3, 32, 75]. This was confirmed in one of our empirical studies, which showed that people who have minority preferences expressed higher privacy risk (see Section 4.2). The majority scenario itself can also have two conditions when the social positions of the group members are not equal as described above (“*boss majority*” and “*peer majority*”).

General privacy concern. General privacy concern is a personal trait that represents an individual’s general tendency to worry about information privacy [54]. Several studies have shown that privacy concerns can significantly reduce trust between consumers and the companies, as privacy concerns decrease trust [18, 133].

Personality. Several studies in the field of behavioral sciences analyze the impact of personality on an individual’s general privacy concern perception [5, 52, 65]. The results, however, are not consistent with each other. Personality is generally modeled using the Five Factors Model (FFM), also known as the Big Five or OCEAN. It models individuals’ personalities with five traits: Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism [14]. Bansal et al. [5], 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 impact of Agreeableness and Neuroticism on privacy concerns. In the context of location-based services, Junglas et al. [52], showed significant effects of Agreeableness on privacy but suggested a negative impact (i.e., more agreeable people were less concerned about their privacy).

- **Extraversion** is a personality dimension linked to being warm, sociable and assertive [1, 15]. Extraverts were also reported to have lower information sensitivity concerns, so as to accommodate their higher need to interact [4]. Therefore, extraversion should be negatively related to user privacy concerns [5, 65, 104].
- **Agreeableness** “involves getting along with others in pleasant, satisfying relationships” [101]. Agreeableness emphasizes trust, altruism, compliance and modesty [1]. Agreeable individuals are also less likely to judge others’ actions as potentially harmful when faced with privacy threats. Hence, their tendency to trust and to be less

suspicious of their environment may reduce their privacy concern. Consequently, they may have lower privacy concerns [5, 52, 65, 104].

- **Conscientiousness** is a personality dimension that emphasizes competence, achievement, self-discipline, and dutifulness [1]. Conscientious individuals have more precaution and foresight, are detail-oriented, and investigate various consequences of a decision, as well as better able to identify potential hazards of disclosing private information [4]. So as conscientious individuals tend to be deliberative, give more attention to details and pay close attention to others' actions, they would also manifest greater concern for protecting their privacy [104].
- **Neuroticism** is a personality dimension characterized by anxiety, self-consciousness, and impulsiveness [1]. It is sometimes referred to as emotional instability, or if reversed as emotional stability (e.g., [5]). In the remainder of the dissertation, I will use the term "neuroticism" as it is the most widely used one. A person with a higher level of anxiety and fearfulness should be more nervous about disclosing their personal information and have a greater privacy concern. A significant and positive effect of neuroticism on privacy concern was found in multiple domains [5, 52, 65].
- **Openness** to new experiences relates to an individual's curiosity, intellect, fantasies, ideas, actions, feelings, and values. Individuals scoring high on this personality trait tend to be less conforming to norms and to have untraditional and widespread interests [1]. They were found to show a high level of scientific and artistic creativity, divergent thinking, liberalism, and only little religiosity [52]. Therefore, and compared to others, open individuals have developed a broader and deeper sense of awareness. As a result of such awareness, they are more likely to be sensitive to things that are threatening [52].

2.5 Conclusions

This chapter mainly focused on studying the state-of-the-art of group explanations. I introduced one of the common approaches to generate group recommendations, namely aggregation strategies which were initially my basis for the content of group explanation. As stated above, there are limited works to develop explanations for group recommendations. And especially none of those consider privacy, which is essential when designing explanations in a group context. This thesis aims to extend work on explaining recommendations/decisions to groups, especially by considering group members' perceived privacy risk. To this end, I presented related works on the privacy calculus theory, which examines information disclosure as a trade-off between disclosure benefit and privacy risk, and antecedents for these.

3

Deciding on the Content for Group Explanations

Consider a group of people trying to make a joint decision, for example, a group of colleagues on a business trip deciding where to visit. People often have different preferences, making it difficult to reach a consensus. Various strategies can be used to predict an item suitable for the group, given the preferences of all group members. Each such aggregation-based prediction strategy, however, has its trade-offs. For example, one strategy will ensure that for each prediction/ recommendation, no one is unhappy in the group, but the resulting recommendations may miss items loved by some users. Another strategy makes sure that, over time, everyone gets the items they love, but only when it is their turn to pick. Explanations, for such recommendations, in this context, act as complementary information and help people comprehend why certain items are recommended. Explanations can be regarded as additional information (i.e., in this thesis this information is textual only) that accompanies the recommendations and serves various goals, such as increasing satisfaction (the ease of usability or enjoyment of the used recommender system) [124]. However, there are many types of information to include and many ways to formulate an explanation, and it is not clear which information should be shown in the explanation for a group.

So in this chapter, we start with the first research question, which addresses what information should be conveyed when generating a group explanation to increase group members' satisfaction (RQ1). We describe three experiments that we conducted to this end. Similar to explanations for single-user recommendations, explanations for groups can be designed based on the underlying recommendation algorithm. So our first experiment (see Section 3.1), evaluates different aggregation strategies used to predict/ recommend items to groups i.e., in terms of perceived satisfaction. An advantage of aggregation strategies is that they can be easily explained to the group, which can be a good start for designing group explanations. A disadvantage is that they are simple and do not, for example, evolve based on users' needs, like machine learning approaches. Based on the results, it seems that it does not matter which aggregation strategy we use as long as we explain that strategy. The second experiment (see Section 3.2), evaluates users' perceptions regarding aggregation strategies and their explanations separately (in isolation). This helps to understand to what extent users' satisfaction eval-

uations are attributed to the explanations or simply the aggregation strategies. It seems that explanations containing information about the aggregation strategy do not significantly benefit users (e.g., increase their satisfaction) in simple scenarios like the one we used (i.e., when there is a few recommended items, and a few number of group members). This suggests that explanations indeed should not contain information about the aggregation strategy. However, these results are not enough to claim that explanations in general are not helpful for group recommender systems. More complex scenarios than the ones we studied might involve a more balanced situation between subgroups with different preferences or a greater number of options to choose from; in such cases, an explanation of the approach used might have an impact. Therefore, the third experiment (see Section 3.3), proposes different group explanation styles for more complex scenarios, for example, when a group member did not receive her favorite item and for a great number of candidate items (10 items compared to 3 items in the previous experiment). There appears to be a benefit of modifying the group explanation style (i.e., repairing vs. reassuring) to the variation in user preferences (i.e., when there is group disagreement or agreement on the recommended item). In addition, user comments highlight the need for protecting certain types of information when presenting an explanation to the group, i.e., group members' ratings of items. This suggests that studying the trade-off between privacy (i.e., protecting certain types of information) and transparency (i.e., disclosing group members information) appears to be a more promising direction than explaining aggregation strategies.

This chapter is based on a conference paper, a workshop paper, and a late breaking results paper:

- Shabnam Najafian, Daniel Herzog, Sihang Qiu, Oana Inel, and Nava Tintarev. *You do not decide for me! Evaluating explainable group aggregation strategies for tourism*. In Proceedings of the 31st ACM Conference on Hypertext and Social Media, pages 187–196, 2020.
- Francesco Barile, Shabnam Najafian, Tim Draws, Oana Inel, Alisa Rieger, Rishav Hada, and Nava Tintarev. *Toward benchmarking group explanations: Evaluating the effect of aggregation strategies versus explanation*. 2021.
- Shabnam Najafian and Nava Tintarev. *Generating consensus explanations for group recommendations*. In UMAP Latebreaking results, 2018.

3.1 Experiment 1: Aggregation Strategies as a Basis for Group Explanation

Several approaches in the literature [74, 76, 85] propose aggregation strategies, which combine the individual preferences of all group members and predict an item that is suitable for everyone. Satisfying the whole group however is challenging especially when group members have different preferences. An explanation in such contexts can indicate possible changes of requirements that help reaching consensus in the group. Similar to explanations for single-user recommendations, explanations for groups can be designed based on the underlying recommendation algorithm. *This experiment evaluates different aggregation strategies used to recommend items to groups as a basis for the content of group explanation (i.e., in terms of user satisfaction, fairness, etc.*

As mentioned, different aggregation strategies may be effective for different situations and groups. For example, items with higher average ratings are not good recommendations when the people in the group have very different preferences. We expand on the state-of-the-art by combining existing aggregation strategies to mitigate the disadvantages of aggregation strategies and avail their advantages, we propose two new explainable aggregation strategies. We assess the impact of our proposed aggregation strategies by comparing them to the *Average* and *Dictatorship* strategies as our baseline strategies in an online study.

In summary, in this experiment, we make the following contributions:

- We investigate which of four explainable aggregation strategies help increase user-perceived satisfaction, fairness, and acceptance.
- We make a setup in such a way people feel more engaged with what is actually being recommended, i.e., by obtaining and using the users' actual travel preferences.

Generally, all investigated aggregation strategies performed comparably well in terms of the satisfaction, fairness and acceptance. There is, however, one strategy (Dictatorship strategy) which received lower average ratings and overall decreased user satisfaction.

In addition to the aggregation strategies evaluated in this study, there are other alternative strategies that could be explored. However, there was no empirical evaluation of these methods with people in the groups and no explanation has been designed yet. It needs more exploration in the future.

The contribution of this study is published as a full paper in Proceedings of the 31st ACM Conference on Hypertext and Social Media [85].

3.1.1 Preliminary Definitions: Aggregation Strategies

This experiment evaluates different aggregation strategies used to recommend items to groups particularly the two aggregation strategies that we propose: *Least+* and *Fair+*, comparing with *Average* and *Dictatorship* proposed in [74].

We describe each of these aggregation strategies with examples used in previous literature [74], with individual ratings for ten items (A to J) for a group of three (John, Adam, and Mary). The highest possible rating is 10. The Sum row calculates the final scores for each item. Group List represents the sorted final recommended list.

Least+ (Least Misery + Most Pleasure + Without Misery) The *Least+* strategy prioritizes, and presents first, items that maximize the rating of the happiest person and at the same time minimize the unhappiness of the saddest person within the group. The *Most Pleasure* strategy considers the highest rating in the group as a group preference rating for the item. The *Least Misery* strategy means that the preferences of a group to items are decided by the lowest rating in the group (the least happy member). The *Least Misery* strategy is one of the prevalent ones and it has been widely applied in traditional group recommender approaches [31]. The *Without Misery* strategy excludes items that anyone in the group rated below a certain threshold. When using the original *Least Misery* and *Without Misery* strategies on their own, items may be selected such that nobody dislikes, but also, nobody really likes. An example of the *Least+* strategy can be seen in Table 3.1. The LM row shows the items' scores after applying the *Least Misery* strategy. This strategy makes a new list of ratings with the minimum of the individual ratings for each item. The next row, MP, shows the items' scores after applying the *Most Pleasure* strategy. This strategy makes a new list of ratings with the maximum of the individual ratings for each item. Finally, the Sum row shows the sum of LM and MP rows. The dashes in the last row indicate that the item will not be considered for recommendations because of the *Without Misery* strategy. So for example, item A which would be one of the top items recommended based on the *Most Pleasure* strategy is excluded from the *Least+* strategy recommendation list, as this item is Adam's least favorite item.

Table 3.1: Applying the *Least+* (*Least Misery* (LM) + *Most Pleasure* (MP) + *Without Misery* (WM)) strategy on an example from [77]. LM and MP rows show the items' scores after applying the LM and the MP strategies respectively. The dashes in the last row show that item will not be considered for recommendations because of the WM strategy.

	A	B	C	D	E	F	G	H	I	J
John	10	4	3	6	10	9	6	8	10	8
Adam	1	9	8	9	7	9	6	9	3	8
Mary	10	5	2	7	9	8	5	6	7	6
LM	1	4	2	6	7	8	5	6	3	6
MP	10	9	8	9	10	9	6	9	10	8
Sum	-	13	-	15	17	17	11	15	-	14

Group List: (E, F), (H, D), J, B, G (threshold 3 out of 10)

Fair+ (Fairness -> Average) The *Fair+* strategy takes turns between users to select their most preferred item, which corresponds to the item with the highest ranking in the rated items list for users. This strategy considers the satisfaction of all the users but could include the most hated item if it is a top item of one member. This strategy in group settings can be characterized as a strategy without favoritism or discrimination towards specific group members [30], compared to *Least+*, where one member could dictate her preferences. In the *Fair+* strategy, one person chooses first, then another, until everyone has made one choice. The next rounds begin with the one who had to choose last in the previous round.

When the rating is the same for multiple items, the item with the higher average rating will be selected. An example can be seen in Table 3.2. In our example, if we start with John first, his favorite items are A, E, or I. We recommend E because it has the highest average. Next, it is Adam’s turn. Adam would like B, D, F, or H. We recommend F because it has the highest average. Mary would choose A (her highest rating). Next, we start with Mary, she would like E, which has already been recommended, and then F, which also has already been recommended. Following the Masthoff [74] approach, we then skip Mary’s preferences in this round and recommend based on Adam’s highest rating. He likes B, D, or H. We recommend H, as that has the highest average. Following this strategy, we could end up with a group list like: E, F, A, H, I, D, B, J, C, G.

Table 3.2: Applying the *Fair+* (*Fairness -> Average*) strategy on an example from [77]. For the sake of readability the sum is not divided by number of group members.

	A	B	C	D	E	F	G	H	I	J
John	10	4	3	6	10	9	6	8	10	8
Adam	1	9	8	9	7	9	6	9	3	8
Mary	10	5	2	7	9	8	5	6	7	6
Sum	21	18	13	22	26	26	17	23	20	22

Group List: E, F, A, H, I, D, B, J, C, G

Average [74] This strategy averages individual ratings and selects items with high average ratings. It does not consider extreme cases, and it is not an optimal method when the individual preferences highly diverge because, for example, extreme low ratings can be balanced out by extreme high ratings. An example can be seen in Table 3.3.

Table 3.3: Applying the *Average* strategy on an example from [77]. For the sake of readability the sum is not divide by the number of group members.

	A	B	C	D	E	F	G	H	I	J
John	10	4	3	6	10	9	6	8	10	8
Adam	1	9	8	9	7	9	6	9	3	8
Mary	10	5	2	7	9	8	5	6	7	6
Sum	21	18	13	22	26	26	17	23	20	22

Group List: (E, F), H, (J, D), A, I, B, G, C

Dictatorship [74] In the Dictatorship strategy (also called ‘Most Respected Person strategy’), only the ratings of one member in the group will be considered for generating the recommendations to the group. In this strategy, the group may be dominated by one person. For example, if you respect highly a person in the group, like your boss, you may all

follow his/her taste. An example can be seen in Table 3.4. In our case we always selected one of the other group members' preferences rather than the active user.

Table 3.4: Applying the *Dictatorship (Most Respected Person)* strategy on an example from [77]. In this example, the ratings of Adam are considered as a dictator.

	A	B	C	D	E	F	G	H	I	J
John	10	4	3	6	10	9	6	8	10	8
Adam	1	9	8	9	7	9	6	9	3	8
Mary	10	5	2	7	9	8	5	6	7	6

Group List: (B, D, F, H), (C, J), E, G, I, A

Summary of strategies and their trade-offs The *Average* and the *Dictatorship* strategies serve as baselines as they have been the most applied strategies by groups (Herzog and Wörndl [43]). Besides, we will evaluate *Least+* and *Fair+* strategies. We believe that it is interesting to compare these two strategies (*Least+* and *Fair+*) as they have complementary strengths and weaknesses. In one (*Least+*), having high average satisfaction by excluding the least preferred item(s) of one or more people. In the other (*Fair+*), having a fair system that might recommend to you your most hated item if it is a top item of another group member (as long as you get to visit the places you really love as well).

3.1.2 Experimental Design

We wanted to understand which of the previously introduced strategies performs better in terms of perceived satisfaction, fairness, and acceptance in high divergence scenarios and low divergence scenarios.

Specifically we aim to address the following research question:

- Which strategy performs the best in which scenario (level of divergence) in terms of user-perceived individual and group satisfaction, perceived fairness, and user acceptance?

For this purpose, we recruited crowd-workers from Amazon Mechanical Turk (MTurk)¹ to conduct a user study.

Preliminaries

To start the user study, we need to create a set of recommendations for groups. For that, we need information for both items to recommend and two synthetic group members (as we have only one active user and need to compose a group of three). For that following steps were needed to be satisfied.

¹<https://www.mturk.com>, retrieved November 2019.

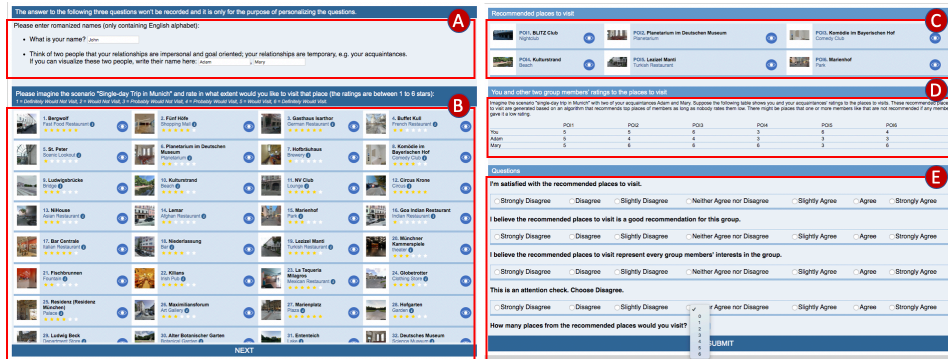


Figure 3.1: Screenshots of the task. First, participants see the left page which includes: (A) input fields to ask participant to enter the names of the imagined group members, (B) the 42 initial POIs to obtain participant’s preferences in the scenario “Single-day Trip in Munich”. Then, participants see the right page which includes: (C) recommended POIs generated by one of the four strategies, (D) the description of the scenario and the explanations of how the strategy works and ratings of the recommended POIs given by the participant and other two group members, and (E) questions for evaluating the recommended POIs.

Data Set

Our first task was to compose a set of recommendations for the user study. We use preferences for different categories from a previous travel-related user study [43]. In that study, every user individually rated all 42 categories (e.g., *Art Museum* and *French Restaurant*), on a scale from 0 (not interested in this category) to 5 (strongly interested in this category). There were 40 groups with 3 members registered for the study. The groups were real, i.e. participants applied as groups and were not randomly assigned. The participants were asked to imagine the scenario “single-day trip in Munich”.

Selecting items to rate.

To obtain the crowd workers’ preferences we wanted to provide them with an initial 42 POIs to rate. We retrieved the most popular POI (in terms of like count) for each selected category (for all 42 categories from the data set) from the social location service Foursquare² as a representative of that category. By using a real data set we increased the likelihood of a realistic rating distribution.

Group composition.

In our experiment, we want to force high divergence and low divergence in the group. To do this, based on crowd workers’ ratings for the 42 initial POIs, we form a group for the crowd worker by picking two synthetic group members from the real tourist data set that we described in the Section 3.1.2. For half of the crowd workers, we select two users with the **highest similarity** compared to the crowd worker and for the other half two users with the **lowest similarity** (dissimilar). We did not consider how similar or dissimilar the other two synthetic group members are to each other, as we were not interested in the average group divergence, but we were only interested in the level of divergence towards

²<https://developer.foursquare.com/>, retrieved April 2019

the real user. A user's travel preferences are represented by a vector of length 42. We used the Pearson's r to determine the similarity/dissimilarity of two user's travel preferences.³

Independent variables

We manipulate the following (independent) variables in this study:

Aggregation Strategies *Least+* and *Fair+* as modified strategies as well as *Average* and *Dictatorship* as baseline strategies.

3

Levels of Divergence In this study we consider two levels of divergence: *high divergence* and *low divergence*. We believe it is more important to study *high divergence* cases because it is more challenging to satisfy all group members when they have different travel preferences. For the sake of comparability we also applied strategies on the groups we predicted to have a low divergence in their preferences and have more similar travel preferences. As explained in Section 3.1.2, we calculated Pearson's r between group members within the group. The range of values for Pearson's r is between -1.0 to 1.0, where -1.0 indicates the strongest negative correlation of travel preferences of two users (contrary preferences) and 1.0 indicates the strongest positive correlation of travel preferences of two users (similar preferences). We consider values between $[-1, 0)$ as *high divergence* and values between $(0, 1]$ as *low divergence*.

Dependent variables

We evaluated each recommended POIs list in terms of four criteria: perceived individual and group satisfaction, perceived fairness and user acceptance. For this purpose, each participant received the following questions on a 7 point Likert scale ranging from 1 (strongly disagree) to 7 (strongly agree):

Perceived individual satisfaction: *"I'm satisfied with the recommended places to visit."*

Perceived group satisfaction: *"I believe the recommended places to visit are good recommendations for this group."*

Perceived fairness: *"I believe the recommended places to visit represent every group members' interests in the group."*

We also asked the users to give us the number of places they would like to visit from the recommended list, *i.e.*, the user acceptance. There is a total number of 6 POIs that we recommend.

User acceptance: *"How many places from the recommended places would you visit?"*

Finally, a free-text comment as the last question was provided for participants to motivate their answers.

³For this purpose we use Pearson's r which measures the linear correlation between two variables and is often used in RSs to identify similar users [64].

Procedure

We designed an online between-subjects experiment in which participants were randomly assigned into a 4 strategies (*Least+*, *Fair+*, *Average*, and *Dictatorship*) \times 2 levels of divergence (low vs high) design. We created $8 = 4 \times 2$ different variations, manipulating strategies and levels of divergence, and each participant only sees one variation. This allows us to evaluate the aggregation strategies in terms of perceived individual and group satisfaction, fairness, and user acceptance in two different scenarios (low divergence vs high divergence). The measurement is the same for all 8 variations as defined in Section 3.1.2.

The experiment consisted of three steps (see Figure 3.1):

Step 1: Some of the participants' individual details were collected, such as demographics (age, gender), education level, and frequency of using apps for recommending touristic places.

Step 2: Next, they were asked to imagine the scenario "single-day trip in Munich" and rate the 42 predefined POIs (see Section 3.1.2). The POIs were augmented by Google Street View for a more accurate preference rating (see Figure 3.1 (B)).

Step 3: According to a participant's initial ratings, for half of the crowd workers a high divergence group and for the other half a low divergence group is created (see Section 3.1.2). Next, a list of POIs is generated, based on all group members' preferences by applying randomly one of the four aggregation strategies, namely: *Least+* (we apply threshold 2 out of 6 in our experiment), *Fair+*, *Average*, and *Dictatorship* (we recall here that for the *Dictatorship* strategy, we always selected one of the other group members' preferences rather than the active user). We presented the top 6 POIs from the generated recommendations since it was a more realistic length for a one-day touristic visit. The recommended POIs were presented as a set, and participants were told that, "this set does not contain any order and can be consumed in any order".

They were presented with explanations of how the strategy works and what the other two group members' ratings are. Figure 3.1 (C, D), illustrates the recommended POIs and their corresponding explanation. Participants can additionally explore each POI using Google Street View.

We asked participants to answer a set of survey questions related to evaluating the recommended POIs in terms of perceived individual and group satisfaction, perceived fairness, and user acceptance (see Section 3.1.2). We also included the following attention check question: "This is an attention check. Choose Disagree" to exclude spurious responses. At the end of the survey, participants were asked to express their opinions regarding the recommended POIs and recommendation algorithm in an open-ended question (see Figure 3.1 (E)).

Hypotheses

In the following, we refer to perceived fairness, acceptance, individual satisfaction, and group satisfaction as F-A-IS-GS to avoid making hypotheses long and hard to differentiate.

Given the trade-off between the strategies and the level of divergence, we hypothesize that:

- **H1)** F-A-IS-GS vary across different aggregation strategies.
- **H2)** F-A-IS-GS differ for different levels of group preference divergence.
- **H3)** F-A-IS-GS vary across different aggregation strategies and levels of group preference divergence (interaction effect).

Statistical analyses

We wanted to determine if there are non-random associations (with regard to the dependent variables) between eight categorical variables (in our case four strategies and two levels of divergence).

To test our hypotheses, we applied the Two-way MANOVA test for between-subjects. Bonferroni correction was applied when multiple tests were conducted. The required sample size was estimated to be 200 participants, this was based on the G*Power analysis for the Two-way MANOVA for between-subjects user studies [26].

3.1.3 Results

In this section, we describe the results of evaluating four different preference aggregation strategies introduced in Section 3.1.1 (*Least+*, *Fair+*, *Average*, *Dictatorship*) to recommend POIs in the context of groups in two scenarios, namely high divergence scenarios and low divergence scenarios.

Participants We recruited 226 participants from MTurk in December 2019. All participants are based in the United States and have overall HIT approval rates of at least 95%. Knowing Munich was not a pre-requisite for the study. Each participant received \$2 as compensation for their time (on average it took 21 minutes of their time to complete the entire task). We excluded 26 participants who failed the attention check question from our data analysis. This resulted in 200 participants (50 per strategy): 38.5% female and 61.5% male. The highest level of education that they held was 29% a high school diploma or equivalent degree, 57% a bachelor's, and 29% a master's degree or higher. Among those, 32.5% use tourism applications (such as Yelp, or Foursquare) less than once a month, 29% at least every month, 22.5% every week and only 8% never used one.

Table 3.5: MANOVA: Wilks Test – it tests the main effect between the strategies, between the levels of divergence and the interaction between the strategies and the levels of divergence on the combined dependent variables

Cases	df	Approx. F	Wilks' Λ	Num df	Den df	p
(Intercept)	1	2255.045	0.022	4	203.000	< .001
strategy	3	2.388	0.872	12	537.379	0.005
divergence	1	2.097	0.960	4	203.000	0.083
strategy * divergence	3	0.558	0.968	12	537.379	0.876
Residuals	200					

Table 3.6: The table shows the average and deviations of the study results for the user perceived individual satisfaction, group satisfaction, fairness, and acceptance. The three first variables are evaluated by 7 point Likert scale, and acceptance shows how many POIs among 6 recommended POIs user would accept. The maximum values for the four measured variables are in bold.

Strategy	Divergence	Individual Satisfaction		Group Satisfaction		Fairness		Acceptance	
		Mean (out of 7)	Std	Mean (out of 7)	Std	Mean (out of 7)	Std	Mean (out of 6)	Std
Least+	Low	5.73	1.07	5.69	0.85	5.34	1.01	2.85	0.46
	High	5.66	1.23	5.46	1.39	5.15	1.61	2.62	0.67
	Total	5.69	1.15	5.58	1.16	5.25	1.32	2.73	0.59
Fair+	Low	5.62	1.23	5.76	1.21	5.45	1.27	2.92	0.39
	High	5.48	1.32	5.73	1.00	5.27	1.21	2.76	0.63
	Total	5.55	1.27	5.75	1.10	5.36	1.23	2.84	0.53
Ave	Low	5.68	0.85	5.78	1.12	5.52	0.79	2.92	0.27
	High	5.61	1.30	5.44	1.00	5.28	0.98	2.91	0.28
	Total	5.65	1.08	5.60	1.06	5.40	0.89	2.92	0.27
Dict	Low	4.97	1.47	5.10	1.11	5.10	1.32	2.83	0.37
	High	4.62	1.72	4.88	1.53	4.85	1.82	2.69	0.55
	Total	4.80	1.59	5.11	1.33	4.98	1.57	2.77	0.47
Total	Low	5.48	1.22	5.53	1.36	5.29	1.29	2.74	0.57
	High	5.35	1.45	5.48	1.00	5.20	1.29	2.88	0.38
	Total	5.41	1.34	5.50	1.19	5.24	1.29	2.81	0.49

H1: F-A-IS-GS vary across different strategies

There was a statistically significant main effect between strategies on the combined dependent variables ($F = 2.390$, $p = .005$; Wilks' $\Lambda = .871$) (see Table 3.5). We did between-subjects effects test (ANOVA) to investigate further the effect on each dependent variable. Tests of the four hypotheses (four variables) were conducted using Bonferroni adjusted alpha levels of .0125 per test (.05/4). Following we discuss its results.

H1.1: in terms of perceived individual satisfaction (IS). The results showed there is a significant difference in perception of individual satisfaction between strategies ($p < .0125$). The post hoc test showed the three strategies (*Least+*, *Fair+*, and *Average*) have significantly higher perceived individual satisfaction compared to the *Dictatorship* strategy ($p < .0125$) (see Table 3.6 for the average and standard deviation values).

H1.2: in terms of perceived group satisfaction (GS). The results showed there is a significant difference in perception of group satisfaction between strategies ($p < .0125$). We did post hoc test to see which strategy led to higher group satisfaction. The results showed that the *Fair+* strategy has significantly higher perceived group satisfaction compared to the *Dictatorship* strategy ($p < .0125$) (see Table 3.6 for the average and standard deviation values).

H1.3: in terms of perceived fairness (F). There was no significant difference between the strategies in terms of perceived fairness.

H1.4: in terms of user acceptance (A). There was no significant difference between the strategies in terms of user acceptance.

H2: F-A-IS-GS differ in levels of group preference divergence

There was no statistically significant main effect between different levels of divergence on the combined dependent variables ($F = 2.097$, $p = .083$; Wilks' $\Lambda = .960$). Therefore, all the sub hypotheses for each individual variable are rejected accordingly. This result will be discussed further in Section 3.1.4.

H3: F-A-IS-GS vary across different strategies and levels of group preference divergence

There was no statistically significant interaction effect between strategies and the level of divergence on the combined dependent variables ($F = 0.558$, $p = .87$; Wilks' $\Lambda = .968$). Therefore, accordingly, all the sub hypotheses for each individual variable are rejected. This result will be discussed further in Section 3.1.4.

Post hoc analysis

All aggregation strategies performed well in terms of all defined dependent variables. There is a difference between strategy performance but the difference is smaller than we expected. We investigated which factors contributed to the surprising results. We particularly looked at differentiation between strategies, the number of common items between strategies, and differentiation between levels of divergence.

Lack of differentiation between strategies

We checked whether we captured the weakness of the applied strategies.

The *Least+* strategy did not exclude the most favorite item from the recommended set. The main weakness of the *Least+* strategy is excluding the highly rated item from the recommended set if it is below a certain threshold for another group member. The *Least+* strategy excluded the top most favorite item for only 6/50 participants. This suggests that the weakness of this strategy did occur for some participants, however, this was not a very common occurrence.

The *Fair+* strategy did not include the least favorite item in the recommended set. The *Fair+* strategy has a weakness that it may include a least rated item of a group member if it is a top item of another group member. We checked how often this happened in this study, and it only occurred for 4/50 participants. This suggests that the weakness of this strategy did occur for a few participants, however, this was not a very common occurrence.

The *Average* strategy did not have extreme low ratings or extreme high ratings. The main weaknesses of the *Average* strategy is that it does not consider extreme cases, and it is not an optimal method when the individual preferences highly diverge because, e.g., extremely low ratings can be balanced out by extremely high ratings. As can be seen in Figure 3.5, our study did not contain very high divergence groups.

The Dictatorship strategy recommended items that represent all group members' preferences and not only the preferences of one member. The main weakness of the *Dictatorship* strategy is that it only considers and recommends based on one group member's preferences. Moreover, it is not an optimal method when the individual preferences are highly divergent. In our set-up, always a member other than the active user dictates her preferences. In Figure 3.2 we show the active user's ratings for each POI, recommended by each strategy. The results show that the average ratings of the active user to the recommended POIs were lower for the *Dictatorship* strategy, compared to the average values of the other three strategies. This can also motivate the difference we found between Dictatorship and other three strategies in terms of individual and group satisfaction (both values were lower for the Dictatorship strategy).




Figure 3.2: The initial ratings (range [1,6]) of the active user for the recommended POIs by each strategy.

Number of common items between strategies was high

To understand the similar results between the strategies, we checked to what extent strategies recommended similar or different POIs. In Figure 3.3 we see that three-quarters of the total amount of POIs, namely 30 out of 42, are in common among all four strategies (10 POIs have never been recommended in any strategy and 20 POIs have been recommended in all 4 strategies). Figure 3.4 shows the pairwise comparison of occurrence of each POI between each pair of strategies. It can be seen that the median of differences for all pairs of strategies is below three (which is a small number). It shows that, in general, strategies did not recommend very different POIs. This might have caused the little difference we found between strategies in terms of the defined variables.

Limited differentiation between levels of divergence

We did not find a significant effect of scenario (high and low divergence groups) in terms of users' perceived satisfaction, fairness, and acceptance. We investigated whether this was due to a lack of differentiation between these groups.

Figure 3.5 shows the box plot of Pearson's r values of a user-user pair in a group (two values per group, active user vs acquaintance 1 and active user vs acquaintance 2).




Figure 3.3: Number of POIs recommended by none, one, two, three or all four strategies.




Figure 3.4: Pairwise comparison of frequency of occurrence of all 42 predefined POIs between strategies: Least+ vs Fair+, Least+ vs Average, Least+ vs Dictatorship, Fair+ vs Average, Fair+ vs Dictatorship, and Average vs Dictatorship.

As can be seen in the Figure, there is a differentiation between high and low divergence groups. However, the median correlation for groups with high divergence is not very low (around -0.12), meaning that, overall, we did not have very divergence groups.

Rather than ensuring divergence on the candidate POIs in the user profiles, we could have enforced diversity in the resulting recommendation list. For example, using a metric such as Intra List Distance (ILD) [119, 139].

However, it can be the case that having a very high divergence is an artificial scenario, especially in the tourism domain where recommended POI are often popular (and liked by many group members). We could have used completely synthetic data where we fixed the divergence levels, but this would decrease ecological validity.

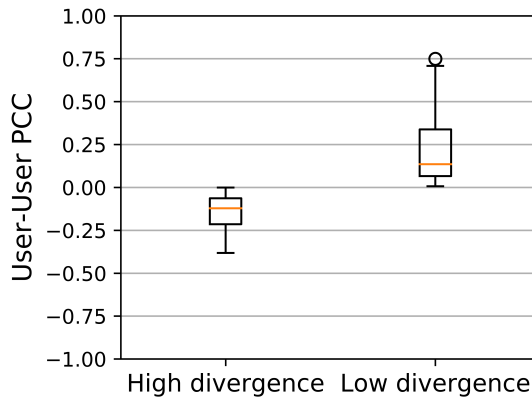


Figure 3.5: The distribution of similarity of travel preferences between users (active user and two non-active users) in high divergence groups and low divergence groups. The similarity is computed using Pearson's r . In our experiments, the similarity of travel preference for high divergence groups is between $[-0.40, 0]$ and for low divergence groups is between $[0, 0.75]$.

3.1.4 Discussion

Despite the differences in strategies used (the *Least+*, the *Fair+*, and the *Average*), both the perceived individual and group satisfaction is comparable and high for all three strategies. However, these three strategies have significantly higher individual and group satisfaction compared to the *Dictatorship* strategy. Some user comments indicated that they are not fully satisfied with the recommended set because it only considered one person's preferences (boss's preferences) in the group.

This can be interpreted as follows: people are more sensitive when the strategy represents the preferences of one member in the group and does not consider other members.

Given the similar results for the strategies and scenarios, we identified in post-hoc analysis that the following factors may have contributed to our results.

Lack of differentiation between strategies. Given that the weaknesses of the different strategies did not happen often, they were all given high ratings.

Number of common items between strategies was high. Given that the strategies often recommended the same POIs, it might be not so important which strategy one applies,

although this is likely to depend on rating distributions, domain, among others.

Lack of differentiation between two levels of divergence. While the two levels of divergence were distinct, this difference could have been stronger.

Limitations

The online experiment used real POIs and user data, which allowed us to have a more realistic scenario, especially by using street view in a European city. On the other hand, using previous ratings constrained our ability to emphasize the differences in scenarios or strategies as might be done in a controlled (but possibly implausible) setting.

Our study only measured the evaluation of one member in a group, when the dictator was someone else in the group, and someone they respected.

This experiment evaluated different aggregation strategies used to generate group recommendations. The results suggested that we probably do not need to consider which strategy we explain as long as we explain the used strategy. So to understand whether the explanation is necessary or not, in the next section (see Section 3.2), we evaluate various aggregation strategies and their explanations in isolation.

3.2 Experiment 2: Aggregation Strategies versus Explanations

In the first experiment, we evaluated underlying methods for deciding what we can recommend to groups, and it seems it does not matter which strategy to use as long as it is not a Dictatorship strategy. The question is, is there still a benefit to explaining those methods. So one way to check if we need an explanation in the first place is to separate the strategy from the output of natural language explanation. This leads to our second experiment, based on the work proposed by Tran et al. [131] — it is unclear to what extent users' fairness perception, consensus perception, and satisfaction evaluations are attributed to the **a) explanations** or **b) simply the aggregation strategies**. Besides, while we agree that the field of group recommendation explanations is a young one, there is no precedent of replication studies, let alone benchmarks and baselines to compare explanations against. The challenges the replication crises have posed in the social sciences and medicine suggest that similar difficulties would be present in other fields involving user studies [96, 97].

This study allows us to address previous limitations in the following ways. For example, 1) we evaluate the *basic explanations* to reproduce the study by Tran et al. [131], with different conditions, such as a different pool of participants and a different setup of the study, which is not affected by participants' learning effects. 2) We can start benchmarking group explanations by comparing this explanation against two other conditions. 3) Finally, the condition without explanations, *no explanation*, allows us to investigate whether the **a) explanations** or **b) simply the aggregation strategies** affect explanation effectiveness (with regard to users' fairness perception, consensus perception, and satisfaction).

To this end, we conducted a preregistered between-subjects user study with 400 participants, where each participant evaluated one aggregation strategy and one explanation type in terms of perceived fairness, perceived consensus, and satisfaction regarding the group recommendations.⁴ In addition, we also test for interaction effects between aggregation strategies and explanation types.

Study participants recognized the MPL strategy as the worst aggregation strategy in terms of perceived fairness, consensus, and satisfaction, while the LMS is recognized as the fairest. Furthermore, we will show that the presence of an explanation, whether basic or detailed, has no significant effect on participants' assessments, either as a main effect or as an interaction effect with the aggregation strategy. These findings suggest that participants are able to discriminate between the different strategies without the need for an explanation and that their perceptions of fairness, consensus, and satisfaction are primarily formed based on the aggregation strategy used to generate the group recommendation. Therefore, we make the following key contributions:

- we reproduce part of the work by Tran et al. [131] with different study conditions and participants pool to understand to what extent their conclusions can be generalized;
- We propose a detailed explanation that describes in great detail the aggregation strategies and provides a reason of why the item was recommended;

⁴To preregister our study, we publicly determined our hypotheses, experimental setup, and data analysis plan before any data collection. The (time-stamped) preregistration can be found at https://osf.io/ghbsq/?view_only=2ef1641bd4204dedbb1f07ddef2ba704 (anonymized for blind peer-review).

- We conduct a preregistered between-subjects user study (N = 400) to investigate the perceived fairness, perceived consensus, and satisfaction of users regarding group recommendations using aggregation strategies, in the presence of two types of explanations (*basic explanations* and *detailed explanations*) and without explanations, *no explanations*;
- We study the main effects of aggregation strategies and explanation types (basic and detailed), as well as interaction effects between them;
- We discuss the challenges of reproducing user studies in recommender systems and highlight design decisions that could influence the process of benchmarking group recommender systems explanations.

All material for analyzing our results and replicating our user study (*i.e.*, document with preregistration of all the hypotheses tested, user study materials, data gathered in the user study, and the analysis scripts) is publicly available.⁵

The contribution of this study is published in Workshop on Perspectives on the Evaluation of Recommender Systems with the 15th ACM Conference on Recommender Systems [7].

3.2.1 Experimental Design

In this study, we aim to address the following research questions:

1. Are there differences between *aggregation strategies* in group recommendation settings regarding users' fairness perception, consensus perception, or satisfaction?

The aim of *research question 1* is to investigate whether the aggregation strategies alone – thus not considering the factor of explanations – can lead to different levels of fairness perception, consensus perception, and satisfaction in users concerning the group recommendations. Previous studies [131] found that an ADD-based explanation was most effective, so we hypothesize that there is a difference between aggregation strategies also when *no explanation* is provided.

2. Do explanations that are based on the group recommendation aggregation strategy at hand increase users' fairness perception, consensus perception, or satisfaction?

We study *research question 2* to investigate whether *explaining* the aggregation strategy at hand to users leads to different levels of users' fairness perception, consensus perception, and satisfaction concerning the group recommendation. In line with the conclusions drawn by Tran et al. [131], we hypothesize that explanations increase each of these variables.

3. Does the effectiveness of explanations (w.r.t. users' fairness perception, consensus perception, or satisfaction) vary depending on the aggregation strategies at hand?

⁵https://osf.io/5xbgf/?view_only=776b756b0f474d35b349d40b49afc355 (anonymized link for blind peer-review.)

Table 3.7: Ratings of group members for the restaurants (1: the worst, 5: the best) from Tran et al. [131].

	Alex	Anna	Sam	Leo
<i>Rest A</i>	2	2	4	4
<i>Rest B</i>	1	4	4	4
<i>Rest C</i>	5	1	1	1

Tran et al. [131] found that explaining the ADD aggregation strategy had the strongest effect on users' perceived fairness, consensus, and satisfaction. However, we study *research question 3* because ADD produces the same recommendation as other strategies (i.e., APP and MAJ – see Section 2.1.1 for definitions and examples of aggregation strategies).

Thus, we expect that there will be no difference between them in the *no explanation* condition. This would suggest that aggregation strategies and explanation types *interact* to affect the perceived fairness, consensus, or satisfaction.

4. Are users' levels of perceived fairness or perceived consensus related to their satisfaction concerning the group recommendations?

We study *research question 4* as a direct replication of findings in Tran et al. [131].

We conducted an online between-subjects user study to investigate our four research questions and test our hypotheses. We presented users with scenarios that reflected one of five different aggregation strategies for group recommender systems, including either no explanation or one of two different explanation types. Evaluating explanation effectiveness (i.e., measuring *fairness perception*, *consensus perception*, and *satisfaction*), we investigated differences between aggregation strategies (*research question 1*) and explanation types (*research question 2*), interactions between aggregation strategies and explanation types (*research question 3*), and relationships between the explanation effectiveness measures (*research question 4*). This section outlines the materials, variables, procedure, participant sample, and statistical analyses related to our user study.

Aggregation Strategies

Our study considered five different aggregation strategies for group recommender systems that have been evaluated in prior work [131]. Each of these strategies aggregates the preferences of several users to obtain a recommendation for the group as a whole [114]. Differently than in [131], we do not consider the *Fairness* aggregation strategy because the explanation types that we propose can not be generated for this strategy. Each aggregation strategy is applied to the rating scenario presented in Table 3.7, where each item (i.e., the three restaurants, Rest A, Rest B, and Rest C) is rated on a 5-star rating scale (i.e., 1 - the worst and 5 - the best). Specifically, we consider the following aggregation strategies, from section 2.1.1: *Additive Utilitarian* (ADD); *Approval Voting* (APP) considering a threshold equal to 3, as in [131]; *Least Misery* (LMS²); *Majority* (MAJ); *Most Pleasure* (MPL).

Explanations

Each explanation is presented after showing the scenario in Table 3.7 and a recommendation generated with one of the aggregation strategies considered (see Section 3.2.1 for more details). We evaluate three types of explanations (see Table 3.8):

Table 3.8: All possible explanation scenarios that participants saw in our study. The explanations describe a restaurant recommendation scenario that participants were exposed to (based on the scenario defined in Table 3.7).

Strategy	No explanation	Basic explanation	Detailed explanation
ADD	<i>“Restaurant B has been recommended to the group.”</i>	<i>“Restaurant B has been recommended to the group since it achieves the highest total rating.”</i>	<i>“Restaurant B has been recommended to the group since it achieves the highest total rating (as the sum of the ratings of all members for Restaurant A is 13 which is higher than other items).”</i>
APP	<i>“Restaurant B has been recommended to the group.”</i>	<i>“Restaurant B has been recommended to the group since it achieves the highest number of ratings which are above 3.”</i>	<i>“Restaurant B has been recommended to the group since it achieves the highest number of ratings which are above a threshold (as the three group members Anna, Sam, and Leo gave it ratings higher than 3).”</i>
LMS	<i>“Restaurant A has been recommended to the group.”</i>	<i>“Restaurant A has been recommended to the group since no group members has a real problem with it.”</i>	<i>“Restaurant A has been recommended to the group since no group members has a real problem with it (as Alex and Anna gave it a rating of 2 which is the highest rating among the lowest ratings per restaurant).”</i>
MAJ	<i>“Restaurant B has been recommended to the group.”</i>	<i>“Restaurant B has been recommended to the group since most group members like it.”</i>	<i>“Restaurant B has been recommended to the group since most group members like it (as 3 out of 4 group members gave it a high rating).”</i>
MPL	<i>“Restaurant C has been recommended to the group.”</i>	<i>“Restaurant C has been recommended to the group since it achieves the highest of all individual group members’ ratings.”</i>	<i>“Restaurant C has been recommended to the group since it achieves the highest of all individual group members’ ratings (as Alex gave it the rating 5, which is the highest rating among all items’ high ratings).”</i>

- *Basic explanations* explain the aggregation strategy at hand. These explanations have been adopted from work published by Tran et al. [131] and refer to *Type 1*.
- *Detailed explanations* explain the aggregation strategy in greater detail by describing the specific reason why a given item has been recommended.
- Additionally, we included a condition *no explanation*, where the aggregation strategy is applied, but no explanation is given. Participants did, however, see the ratings of the other group members in this condition.

Variables

Independent Variables

- **Aggregation strategy** (categorical, between-subjects). Each participant was exposed to a scenario that reflected one of the five aggregation strategies (i.e., ADD, APP, LMS, MAJ, or MPL; see Section 3.2.1).
- **Explanation type** (categorical, between-subjects). Each participant saw either *no explanation*, a *basic explanation*, or a *detailed explanation* (see Section 3.2.1).

Dependent Variables

We measured each of our three dependent variables by asking participants to respond to a statement on a seven-point Likert scale ranging from “strongly agree” (scored as -3) to “strongly disagree” (scored as 3).

- **Perceived fairness** (ordinal): “The group recommendation is fair to all group members.”
- **Consensus** (ordinal): “The group members will agree on the group recommendation.”
- **Satisfaction** (ordinal): “The group members will be satisfied with regard to the group recommendation.”

Descriptive Variables

In addition to the independent and dependent variables that we used for hypothesis testing, we collected data on two different descriptive variables for the demographic description of our sample. Participants could also select a “prefer not to say” option for these variables.

- **Age** (categorical). Participants could select one of the options *18-25*, *26-35*, *36-45*, *46-55*, *>55*.
- **Gender** (categorical). Participants could select one of the options *female*, *male*, or *other*.

Procedure

Our study consisted of two subsequent steps. We introduced participants to the study during the first step (after participants had agreed to informed consent) and asked them for their gender and age. The second step of our study started with the following scenario (taken from Tran et al. [131]):

“Assume there is a group of four friends (Alex, Anna, Sam, and Leo). Every month, a group decision is made by these friends to decide on a restaurant to have dinner together. To select a restaurant for dinner next month, the group has to make the same decision again. In this decision, each group member explicitly rated three restaurants (Rest A, Rest B, and Rest C) using a 5-star rating scale (1: the worst, 5: the best). The ratings given by group members are shown in the table below:”

After that, Table 3.7 is shown. Subsequently, participants saw a group recommendation and (or without) explanation depending on which aggregation strategy and explanation type they had been assigned to (see Table 3.8). We then measured perceived fairness, perceived consensus, and satisfaction (see Section 3.2.1). We also included an attention check where we specifically instructed participants on what option to select. Finally, participants have the option to explain their answers to the three items in an open text field. Our study had been approved by the ethics committee of our institution and pre-registered prior to any data collection.

Participants

Before data collection, we computed the required sample size for our study in a power analysis for a between-subjects ANOVA (Fixed effects, special, main effects, and interactions; see Section 3.2.1) using the software *G*Power* [27]. Here, we specified the default effect size $f = 0.25$, a significance threshold $\alpha = \frac{0.05}{11} = 0.005$ (due to testing multiple hypotheses; see 3.2.1), a power of $(1 - \beta) = 0.8$, and that we will test $5 \times 3 = 15$ groups (i.e., 5 different aggregation strategies for 3 different explanation scenarios). We performed this computation for each of our hypotheses using their respective degrees of freedom. This resulted in a total required sample size of at least 378 participants.

We thus recruited 400 participants from the online participant pool *Prolific*⁶, all of whom were proficient English speakers above 18 years of age. To maintain high-quality answers, we selected only participants with an approval rate of at least 90% and who participated in at least ten prior studies. Each participant was allowed to participate in our study only once and received £0.63 as a reward for participation. We excluded one participant from data analysis because they did not pass the attention check we included in the experiment. The resulting sample of 399 participants was composed of 61% female, 38% male, and 1% other participants. They represented a diverse range of age groups: 28% were between 18 and 25, 29% between 26 and 35, 17% between 36 and 45, 14% between 46 and 55, and 13% were above 55 years of age. We randomly distributed participants over the 15 conditions (i.e., exposing them to one of the five aggregation strategies and one of the three explanation types).

⁶<https://prolific.co>




Figure 3.6: Participants’ mean *fairness perception*, *consensus perception*, and *satisfaction* across explanation types on scales from -3 (“strongly disagree”) to 3 (“strongly agree”; see Section 3.2.1). Colors indicate aggregation strategies: Additive Utilitarian (ADD), Approval Voting (APP), Least Misery (LMS), Majority (MAJ), and Most Pleasure (MPL). Error bars represent the standard error of the mean.

Statistical Analyses

For each of the three dependent variables in our study (i.e., *fairness perception*, *consensus perception*, and *satisfaction*), we conducted a two-way analysis of variance (ANOVA) using *aggregation strategy* and *explanation type* as between-subjects factors. These three ANOVAs were used to test a total of nine hypotheses (i.e., **H1a** – **H3c**). Specifically, each of them tested main effects of *aggregation strategy* (**H1a** – **H1c**) and *explanation type* (**H2a** – **H2c**) as well as the interaction between these two variables in affecting the dependent variables (**H3a** – **H3c**). We chose this type of analysis despite the anticipation that our data may not be normally distributed (i.e., violating an ANOVA assumption) because ANOVAs are usually robust to Likert-type ordinal data [95]. We additionally performed two Spearman correlation analyses to test hypotheses **H4a** and **H4b**. Because we thus tested 11 different hypotheses, we did not handle the typical significance threshold of 0.05. Applying a Bonferroni correction [89], we lowered the significance threshold to $\alpha = \frac{0.05}{11} = 0.0046$.

Because we found significant main effects related to our first six hypotheses (**H1a** – **H2c**; see Section 3.2.2), we conducted Tukey posthoc analyses to investigate specific differences between the aggregation strategies and explanation types. The *p*-values from this analysis are automatically adjusted to correct for multiple testing (i.e., written as p_{adj}).

3.2.2 Results

In this section, we present descriptive statistics as well as the results of the hypothesis tests outlined in Section 3.2.1.

Descriptive Statistics

Participants’ distribution over the 15 different conditions (i.e., all possible combinations between the five aggregation strategies and the three explanation types) was balanced: each condition was shown to 6% to 7% of participants. On average, participants spent 2.9 (sd = 2.2; no notable difference between conditions here) minutes on the task. Qualitative feedback from participants suggested that the scenario and task were understandable. Participants had a slight overall tendency to perceive fairness, consensus, and satisfaction

Table 3.9: Results of three two-way ANOVAs for the dependent variables (DVs) *fairness perception* (left), *consensus perception* (center), and *satisfaction* (right). Per effect, we report the F -statistic, p -value, and η_p^2 effect size. The terms “aggr” and “expl” represent the independent variables *aggregation strategy* and *explanation type*. Colons indicate interaction effects, and asterisks statistical significance.

DV: Fairness Perception				DV: Consensus Perception				DV: Satisfaction			
	F	p	η_p^2		F	p	η_p^2		F	p	η_p^2
(H1a) aggr	36.19	<0.001*	0.27	(H1b) aggr	38.89	<0.001*	0.29	(H1c) aggr	49.57	<0.001*	0.34
(H2a) expl	0.35	0.71	0.00	(H2b) expl	0.14	0.87	0.00	(H2c) expl	0.15	0.86	0.00
(H3a) aggr:expl	0.68	0.71	0.01	(H3b) aggr:expl	0.75	0.65	0.02	(H3c) aggr:expl	1.25	0.27	0.03

in the scenarios, as 51%, 51%, and 56% at least somewhat agreed on these three items, respectively. Figure 3.6 shows participants’ mean *fairness perception*, *consensus perception*, and *satisfaction* across explanation types and split by aggregation strategies.

Hypothesis Tests

Research question 1. We investigated the differences between aggregation strategies regarding explanation effectiveness. We found significant differences between the five aggregation strategies concerning all three dependent variables *fairness perception*, *consensus perception*, and *satisfaction* (H1a – H1c; $F = [36.19, 49.57]$, all $p < 0.001$, $\eta_p^2 = [0.27, 0.34]$; see Table 3.9). So, overall, participants expressed different levels regarding these three variables based on which aggregation strategy they were exposed to. Posthoc analyses revealed that MPL led to lower levels on all three variables compared to all other aggregation strategies (all $p_{\text{adj}} < 0.001$). The only other significant differences we found between aggregation strategies was that APP ($p_{\text{adj}} = 0.004$) and MAJ ($p_{\text{adj}} = 0.005$) each led to lower fairness perception compared to LMS. In sum, participants – irrespective of which explanation type they saw – viewed MPL as significantly less fair, consensual, and satisfying compared to other strategies and judged MAJ as well as APP as less fair compared to LMS.

Research question 2. We studied differences between explanation types (i.e., *no explanation*, *basic explanation*, or *detailed explanation*). We found no significant differences between the three explanation types regarding all three dependent variables (H2a – H2c; $F = [0.14, 0.35]$, $p = [0.71, 0.86]$, all $\eta_p^2 = 0.00$; see Table 3.9). So, our results contain no evidence for a difference between explanation types concerning our three dependent variables.

Research question 3. We researched interactions between aggregation strategies and explanation types regarding explanation effectiveness. There were no significant interaction effects between the five aggregation strategies and the three explanation types (H3a – H3c; $F = [0.65, 1.25]$, $p = [0.27, 0.71]$, $\eta_p^2 = [0.01, 0.03]$; see Table 3.9). The effect of explanation types on participants’ *fairness perception*, *consensus perception*, and *satisfaction* did not differ by applied aggregation strategy.

Research question 4. We studied associations between explanation effectiveness measures. In line with the findings of Tran et al. [131], Spearman correlation analyses revealed

significant positive relationships between fairness perception and satisfaction ($\rho = 0.71$, $p < 0.001$) as well as between consensus perception and satisfaction ($\rho = 0.76$, $p < 0.001$). This means that, as participants' fairness and consensus perception increased, satisfaction also increased.

3.2.3 Discussion

In the following, we look closer at our results and their implications. We discuss the difference between aggregation strategies, the difference between different explanation levels, and the effect of the chosen scenario (a group with heterogenous ratings). We conclude with lessons learned for future benchmarking studies in explanations research and limitations of our study.

The Differences Between Aggregation Strategies

As shown in Section 3.2.2, there are differences between the aggregation strategies in terms of perceived fairness, consensus, and satisfaction. The MPL strategy appears to obtain the lowest scores, regardless of the type of explanation received. Furthermore, MAJ and APP are perceived to be less fair than LMS. We discuss how these results may have interacted with the presented scenario in Section 3.2.3. However, these results are in contrast with the findings of Tran et al. [131], where the same scenario was used. In such work, the MAJ and ADD strategies scored better than the LMS strategy. An explanation of this difference can be the different design of our experiment: we implemented a between-subject design to guarantee the independence between the conditions; on the contrary, in [131], each user evaluated six strategies and was exposed to different explanation types. Although the strategies were presented in a randomized order to reduce biases, it is possible that the user used an explanation type seen first as a reference point to compare with in the following evaluations. Even if this should not have had an impact on a specific aggregation strategy, it could have introduced noise in the users' evaluations. Since we also had a *no explanation* condition, we asked participants to evaluate the recommendation. In contrast, Tran et al. [131] asked the participants to evaluate the explanation provided. This allows us to evaluate the effect of the aggregation strategy separately from the explanation, while in previous work, the evaluation of the explanation was influenced by the evaluation of the aggregation strategy.

The Role of Explanations

The results presented showed no significant difference between the different types of explanations. Furthermore, no interaction effects between the explanations and the aggregation strategies with regard to the measured dependent variables (perceived fairness, consensus, and satisfaction) were found. However, these results are not enough to claim that the explanations are not useful for group recommender systems. First, it must be considered that the scenario used was particularly simple to evaluate. More complex scenarios might involve a more balanced situation between subgroups with different preferences or a greater number of options to choose from: such factors might complicate the assessment; in such cases, an explanation of the approach used might have an impact.

Moreover, the strategies presented here represent baselines for group recommenders. Therefore, it is necessary to formalise the explanations for these strategies, as they serve

as a reference against which more articulated strategies can be compared. But the most recent lines of research in group recommenders try to integrate in the recommendation generation process personal factors (experience in the domain [33] or personality [93, 108, 112]) as well as social factors (tie strength [6], centrality of group members in the group social network [20], group diversity [21]). In such cases, an explanation may have an impact on the transparency and comprehensibility of the system and result in different evaluations regarding fairness perception, consensus perception, and satisfaction. This of course, also leads to privacy issues concerning which personal information of one or more individuals can be mentioned in an explanation.

3

The Link Between Fairness, Consensus, and Satisfaction

The correlation between fairness perception (or consensus perception) and satisfaction, already reported in Tran et al. [131], and also shown in our results, confirms the close connection between these concepts. A solution perceived as less fair is also perceived as less satisfactory, and a less satisfactory solution is unlikely to be accepted by the group. This confirms that these aspects, sometimes considered secondary, are crucial and that a group recommendation system must take them into account, both in the generation of recommendations and in their evaluation.

Lessons Learned for Benchmarking

Here we present lessons learned for benchmarking group recommender systems explanations.

Report on Participant Recruitment. Numerous platforms that can be used to outsource user studies [103], such as Prolific and Amazon Mechanical Turk, to name a few. In addition, user studies could be outsourced to particular users, such as students or staff members. Filtering conditions, such as those for quality control, also affect which demographics take part in a study. More generally, any selection of study participants can influence the outcome of the evaluation, which should not be generalized outside the scope of the scenario [8]. Therefore, we recommend a thorough report on how participants were recruited.

Report Study Design and Statistical Analysis Rigorously. The choice of the quantitative study, between-subjects, within-subjects, or mixed designs is also influencing the conclusions that can be drawn, as well as the statistical analysis that should be applied. In any case, randomizing participants to conditions is of paramount importance, regardless of the study design. More personalized study designs, such as the one conducted by Tran et al. [131], should clearly specify how each scenario has been allocated to participants to be able to replicate them. We, in particular, recommend more rigorous reporting of how randomization is performed, as well as sharing scripts to support replication and comparison.

Ensure consistency in measurement or motivate changes well. In separating the evaluation of explanations and aggregation strategy, we found it was no longer feasible to ask participants to evaluate the explanations rather than the resulting recommendation.

In addition, compared to Tran et al. [131], we ask study participants to rate explanations' effectiveness on a 7-point Likert scale instead of a 5-point Likert scale since this ensures greater robustness in the use of ANOVA analysis, according to [95]. While these changes may not have affected the results, such changes in the design must be described and motivated when attempting to benchmark such user studies.

Report on Completeness. We found that certain aggregation strategies can not be explained in certain instances or scenarios. In this study, this was the case of the *Fairness* strategy, which is well-suited for repeated decisions, but less applicable for single decisions as in our case. We recommend that future work not only describes the cases where explanations can be generated but also describes the edge cases for which they cannot.

Consider the Effect of the Scenario. The proposed scenario in this work was selected to specifically study groups with heterogeneous preferences. However, this choice is likely to have affected our specific results. For example, the MPL strategy in this specific scenario recommends a solution that displeases at most three out of four group members (Rest C, see Table 3.7). It is not surprising, therefore, that it is identified as the least fair, least satisfactory strategy, and the one on which it is most difficult to reach an agreement. The result might have been different if it displeased fewer members. We, therefore, recommend not only clearly reporting the scenario used but discussing its implications.

Consider effects of the role of the participant in a group. The evaluations are given in this study on the basis of an *external* evaluator who may be more unbiased (than someone within the group). A user within the group may be influenced by their own preferences. Furthermore, the assessment of the fairness of a scenario will likely differ depending on whether it favors the user, e.g., if MPL displeases two users and whether the active user is one of them.

Limitations

Here we discuss the limitations of our study.

Recommendations and explanations are not evaluated by group members. As previously mentioned, in line with the evaluation approach in Tran et al. [131], our study participants were asked to evaluate the recommendations as external evaluators. This means that study participants were not members of the group. We hypothesize, however, that their evaluations could be different when part of the group. Deciding for an evaluator that is part of the group would entail controlling more cases, such as when the evaluator is in the majority preference, minority preference, or a tie preference.

We do not measure nor capture the reasoning process of the study participants regarding recommendations. In the condition with *no explanations*, we provide a mere description of the recommendation. However, we do not capture how study participants reflect on the recommendation or to what extent they understand it. Prior literature [35, 59, 134], however, provides several directions for measuring recommendation understandability, which could be investigated in future work. Nevertheless, our descriptive

analysis in Section 3.2.2 shows that participants spent a similar amount of time completing each explanation condition. This could potentially mean that they spent a similar amount of time analyzing their fairness and consensus perception, as well as satisfaction regarding the recommended restaurant. In addition, comments provided by study participants in the condition with *no explanation* indicate that they understood the recommendation well, only by looking at the scenario in Table 3.7. For example, one participant mentions the following regarding the recommendation of MPL: “*Although Alex rated the restaurant as a 5, the others only put 1. This makes it unfair to the majority of the group.*”. Similarly, another participant mentioned that “*It was not the lowest choice among all of the members, and some enjoyed it quite a bit. Ideally, this means they would not mind...*” about the LMS recommendation.

3

Recommendations are provided for unnamed restaurants. We did not want to influence participants’ decisions by providing real restaurant names as recommendations. This helped us control for the potential bias that could have been added while showing a real restaurant name. Such normalization, however, could potentially influence the assessments of the study participants compared to a customized recommendation. Another limitation of our study is that all recommendations are in the restaurants’ domain. Different recommendation domains could be differently perceived in terms of fairness, consensus and satisfaction. In particular, the investment related to the domain considered has shown to have an impact on the evaluation of the recommendations [125]; the restaurant domain is generally perceived as a medium-low investment compared to other domains suitable for group recommendations, such as a shared apartment.

Ethical considerations

Our study received ethical committee approval from the Delft University of Technology. We commit to making all data and code publicly available for the community to be able to replicate and reproduce our study and results. However, the raw results of our user study are anonymized, *i.e.*, we do not publish participants’ identifiable information such as user IDs. In our user study, we aimed for a balanced participants pool in terms of gender and age. While we do observe that participants are more skewed towards younger people, we also note that all age categories are well covered. Furthermore, our participants’ pool is also balanced in terms of gender. Thus, we consider that our results are not biased as an effect of poor age and gender diversity among our study participants. With regard to socio-economical and geographical aspects, we do not have a clear indication. Prior research [103], however, positions the Prolific platform high in terms of socio-economical and geographical diversity.

Our findings suggest that explanations containing information about the aggregation strategy do not significantly benefit users (e.g., increase their satisfaction) in simple scenarios like the one we used in this experiment. In the next section (see Section 3.3), we evaluate various group explanations in more complex scenarios, for example, where there is a group disagreement on the recommended item.

3.3 Experiment 3: Formulating Group Explanations

In the previous experiment, we did not find that natural language explanations made a significant difference in terms of users' fairness, consensus, and satisfaction perception compared to just showing the effects of the aggregation strategy. However, these results are not enough to claim that explanations are never useful for group recommender systems. First, it must be considered that the used scenario was particularly simple to evaluate. More complex scenarios might involve a more balanced situation between subgroups with different preferences or a greater number of options to choose from; in such cases, an explanation of the approach used might have an impact. In this experiment, we suggest different group explanations styles that could be *reassuring* or *repairing* for more complex scenarios, depending on the preferences composition in the group. In addition, the information contained in the explanation is influenced by the scenario. In this scenario, we propose a repairing explanation style for the explanation content, when there is a group disagreement on the recommended item. For instance, *"The system detected you might not like this item, but it is the item Mary prefers most. You made your choice in the previous round; now it's Mary's turn"*. In contrast, we propose a reassuring explanation style for the explanation content, when all group members agree on the recommended item. For instance, *"The system detected that you all would enjoy this item. Moreover, Anna and Bob will love it"*. Participants preferred short, informal, and encouraging explanations. However, when maximal misery (not getting their liked item) was expected, a more complicated explanation was acceptable. Users' comments highlighted the need for privacy when revealing their personal information, like their preferences in the group.

In summary, in this experiment, we make the following contributions:

- We introduce different explanation styles, i.e., repairing vs reassuring based on the applied scenario.
- We introduce different scenarios depending on group members' preferences composition in the group.
- We evaluate explanations resulting from different aggregation strategies regarding user satisfaction. We assessed how four different explanations styles influenced user satisfaction in five different scenarios in structured interviews.

The contribution of this study is published as a LBR paper in the Proceedings of the 26th ACM Conference on User Modeling, Adaptation and Personalization [84].

3.3.1 Preliminary Definitions: Explanation Styles

The choice of the used aggregation strategy also influences the types of explanations we can generate. We give examples of how this might look for different explanation categories in Table 3.10.

To keep a group satisfied during the recommendations we need to consider the preferences of all the people in the group. This can be challenging when the preferences of individual group members diverge. An explanation in such contexts can indicate possible changes of requirements that help improve user satisfaction. In the context of group, such repair-related explanations help group members understand the constraints of other

group members and decide in which way their own requirements or preferences should be adapted [30].

Table 3.10 demonstrates the proposed explanation categories, which we also explain in relation to our study below:

Repairing versus reassuring In this study, we proposed to generate explanations for the *Least+* aggregation strategy on repairing inconsistency category with *pleasure* as a basis, and for the *Fair+* aggregation strategy on the same category with *fairness* as a basis. Both describe group disagreement situations. We call the explanations of this category repair explanations. Here is an example of fairness basis: *"The system detected you might not like song 1 but it is the song Mary prefers most. You made your choice in the previous round, now it's Mary's turn"*.

For comparison, we also study the situation where all group members agree on the selected item. In this study, we call these *reassuring explanations*, which are similar to the positive explanations which have been discussed in Quijano-Sanchez et al. [108] work. For instance, *"The system detected that you all will enjoy this song. Moreover, you and Adam will love it"*.

In our study, we put persuasiveness (as defined by Quijano-Sanchez et al.) under the repair inconsistency category.

Complete and vital The privacy preserving category is used when the underlying recommendations are aggregated models instead of aggregating recommendations for individual users, this approach constructs a group preference model (group profile) that is then used for determining recommendations. The advantage of applying group preference models is that the privacy concerns of users can be diminished [30].

In this study, we represent this as complete explanations and explanations with only vital information. With *complete information* we describe the ratings of everyone in the group, however with *vital information* we only report partial information. More specifically, for the least misery part of the strategy, we report the member of the group with the minimum personal value score for the item, i.e., the member that is responsible for this selection. Similarly, for the most pleasure part of the aggregation strategy, we report the member of the group with the maximum personal value score for the item. Finally, for the fairness strategy, we report with each item the member of the group whose turn it is, i.e., the member direct towards this selection [98]. Following examples are represented as complete and vital information respectively, *"You, Mary and Adam have rated song 5 with values 4, 10, 5 respectively. Song 5 is recommended because it avoids dissatisfaction within the group due to the lowest rating determined for you and supports the highest rating determined for Mary."* and *"Song 5 is recommended because it avoids dissatisfaction within the group due to the lowest rating determined for you and supports the highest rating determined for Mary"*.

3.3.2 Experimental Design

In the previous sections, we introduced aggregation strategies for generating recommendations for groups of users. These naturally influence the explanations that are generated. In addition, whether there is a disagreement in preference will also influence the resulting

Table 3.10: Explanation categories and examples.

Categories	Example
Privacy Preserving	<i>A majority thinks that it is a good choice. Some group members think that it is an excellent choice [30].</i>
Repairing Inconsistency (Persuasiveness)	<i>Although your preference for this item is not very high, your close friend X (who you highly trust) thinks it is a very good choice [108].</i>
Repairing Inconsistency (Fairness)	<i>The interest dimensions favored by user u1 has been given more consideration since u1 was at a disadvantage in previous decisions [30].</i>
Repairing Inconsistency (Pleasure)	<i>Item y is recommended because nobody hates it in the group due to the lowest rating determined for user a and support the the highest rating determined for user b.</i>
Reassuring	<i>Additionally, Jaime, who you trust the most, would really love this movie, so why not give it a try [108].</i>

explanation; in this study, we study *repair* (the group disagrees) and *reassuring* (the group agrees) explanations.

This is a formative and exploratory evaluation with an aim to study how explanations should be designed to maximize satisfaction even when no consensus exists.

Specifically we aim to address the following research question:

- Which explanation performs the best in which scenario in terms of user-perceived satisfaction?

We used the layered evaluation proposed by Paramythis et al.[102], which suggest that for effective adaptation, the process needs to be decomposed and evaluated in layers. This ensured accurate input to the explanation presentation layer. To create a controlled experiment we used synthetic ratings for individual users. The ratings could be potentially the output prediction of any recommendation algorithm, such as Collaborative Filtering, Content-based filtering and so on. If we chose any particular algorithm, the quality of the prediction would affect the quality of the sequence and would affect the quality of the explanation.

Study Design

In a structured interview⁷ participants were asked to assume that they would be listening to a playlist with two of their friends during their travel sitting in a car. Each participant conducted the particular individually (with the interviewer).

⁷<https://goo.gl/DA7Kmf>, retrieved April 2018

They were given a sample of individual ratings (based on synthetic data) for 10 songs just for themselves, not for their friends. They were told that the system has selected a sequence of songs for them and has provided an explanation for the selected sequence.

Next, they were asked how satisfied they are with the presented explanation, what can be made better, or what they liked about that explanation as well as how it affects their satisfaction for the recommended song.

The sequences resulting from the two modified aggregation strategies (*Least+* and *Fair+*, see Section 3.1) were at most 10 songs or less because in some cases strategy resulted in a short sequence. We used "you" when referring to the participant. We explained that their real names would be replaced with their names in the real explanation, and that their real friend's names would be used in the place of "Adam" and "Mary".

3

Procedure

The main user task was to "report her satisfaction degree regarding the proposed explanations in different scenarios". In addition to that she gave her feedback on what can be made better or what she liked about that explanation.

The *independent variables* manipulated in this interview were:

Explanation style: repair or reassuring explanations (2) * only vital information versus complete information (2).

Scenarios: two for each aggregation strategy, & one where all users agree (5).

These were studied in a within-subjects design with each participant seeing all versions. To control for order effects, the scenarios and explanation styles were counterbalanced across participants.

Explanation Categories

We presented four types of explanations:

1. Repair-related explanation with vital information
2. Repair-related explanation with complete information
3. Reassuring explanation with vital information
4. Reassuring explanation with complete information

Scenarios

Users were asked to imagine that they were listening to the playlist with two friends in a car during a roadtrip. The different scenarios studied were:

Scce 1: A song that the user hate has been selected resulting from the *Fair+* aggregation strategy.

Scce 2: The song(s) that the user really likes has not been selected at all resulting from the *Least+* aggregation strategy.

Scce 3: The song(s) that the user really likes has not been selected yet resulting from the *Least+* aggregation strategy.

Scen 4: All group members agree on the selected song (Baseline).

Scen 5: It is the user turn to pick and her favorite song has been selected resulting from the *Fair+* aggregation strategy.

3.3.3 Results

Designing explanations that improved user satisfaction was the goal in this study. We have proposed different types of explanations based on different sequence constructing aggregation strategies, and we investigated user impressions of these explanations in different scenarios. Figure 3.7 summarizes the results by explanation and scenario. The vertical axis shows average satisfaction for each explanation per scenario. Moreover, the error bars indicate the standard deviation (SD) of these results. Due to the small sample size and that this study is exploratory, we have not performed statistical analysis.




Figure 3.7: Average satisfaction, by different explanation types and scenarios. Whisker plot depicts 1 SD.

Participants

Sixteen participants from the staff and student population of Delft University of Technology participated voluntarily in the experiment. They were at least 18 years of age, and 20% female.

Which Explanation Performed Better

Comparing between the aforementioned four types of explanations, explanation 3 (reassuring with vital information) performed better in terms of satisfaction *regardless of the scenario* in which it was presented. The average satisfaction for explanation 3 are ($m = 3.4$, $SD = 1.15$), ($m = 3.6$, $SD = 1.09$), ($m = 3.9$, $SD = 1.06$), ($m = 4.6$, $SD = 0.51$), ($m = 3.8$, $SD = 1.2$), in scenarios 1, 2, 3, 4, and 5 respectively. In particular, explanation 3 in scenario 4 has the highest average satisfaction ($m = 4.6$, $SD = 0.51$). In addition to scoring the explanations, we asked participants why they liked that specific explanation and why not. Some reasons that they liked explanation 3 are as follows: "The explanation is easy to understand", "The encouraging tone.", "Nice, friendly, clear and short", "The explanation is short and concise".

Table 3.11: Possible explanation purposes for groups

Aim	Definition
Transparency	Explain how the system works [123]
Trust	Increase users' confidence in the system [123]
Privacy-Preserving	Preserving users' confidential data, like their preferences
Effectiveness	Help users make good decisions [123]
Persuasiveness	Convince users to try or buy [123]
Group Satisfaction	Increase the average ease of use or enjoyment of all group members
Individual Satisfaction	Increase the ease of use or enjoyment of each member of the group
Single Item Satisfaction	Increase the average ease of use or enjoyment of all group members for each single item
Several Items Satisfaction	Increase the average ease of use or enjoyment of all group members for several items

The traits were mostly mentioned by participants include brevity, simplicity, friendly tone, as well as clear and understandable content.

3.3.4 Discussion

Following this, we look closer at our results and their implications.

Influence of Explanation Category

We compared *vital information* (explanations 1,3) vs *complete information* (explanations 2,4). We found that for *all scenarios, except scenario 2* vital information led to more satisfied participants.

In contrast, the satisfaction for scenario 2 is slightly higher for explanation 2 (compared to explanation 1), with the complete repair explanation. In the case of scenario 2, the increased complexity of the complete information may help users to deal with missing a song they really like. However, both explanations 1 and 2 have low scores. Note that explanation 3 (reassuring-vital) still outperforms explanation 4 (reassuring-complete) also for scenario 2. I.e., this is similar to the other scenarios.

Next, we compared *repair-related explanation* vs *reassuring explanation*, and found that reassuring explanation performed better rather than repair-related explanation. According to the users' feedback we can infer that they preferred to receive positive and encouraging explanations rather than receiving explanations showing misery or dissatisfaction of any of the group members.

Influence of Scenarios

Overall, participants were more satisfied with the explanations in *scenario 4* (Assume all group members like the selected song) and *scenario 5* (Assume it is your turn and you got your favourite song). This can be expected as these are positive scenarios for the users. At the end of each scenario we asked participants "*how the explanation influenced their satisfaction regarding the selected song in general*", results are demonstrated as General in Figure 3.7. Scenario 5 has higher general (with the song) satisfaction than scenario 4 when comparing across all explanation styles ($m=4.31$, $SD=0.8$). This suggest that users care more about their own preferences than global satisfaction in the group.

In scenario 1 (Assume a song that you really hate is now playing) the difference between average satisfaction of explanation 1 and 3 is small with values ($m=3.25$, $SD=1.12$) and ($m=3.4$, $SD=1.15$) respectively. Some comments for explanation 1 in this scenario are "Sad result, but the explanation makes it a bit better.", "It provides proper reasoning as to why the song was selected.", "It acknowledges I don't like the song." or "Seems fair! I'm willing to let them enjoy." and feedback for the explanation 3 include: "It doesn't acknowledge my dislike, but I like the part that my friends will like it." or "It's short and informal."

Although we only asked users about their impressions of the explanations, they also gave feedback regarding the applied aggregation strategy. This was mostly for scenario 1 where we applied the *Fair+* aggregation strategy: "It feels strange that the song is chosen only because it's Mary's favorite song. I would expect a solution where none of the extreme valued songs are chosen to keep the overall satisfaction of both of us higher.", "The songs that anyone likes as little should be kept to the last, even if it is someone's favourite.". The users' feedback illustrate that the *Fair+* aggregation strategy was found to be less satisfying than the *Least+* aggregation strategy.

Influence of Wording

The results suggest that explanation type 1 in scenario 2 has the lowest average satisfaction but with the highest SD ($m=2.75$, $SD=1.34$). This is the explanation: "Song 5 is recommended because it avoids misery within the group due to the lowest rating determined for you and support the highest rating determined for Mary."

This result suggests that although the average satisfaction is low for this explanation participants' opinion vary about that. Some reasons that participants mentioned are "The word "misery" is too strong." or "The explanation sounds a bit complicated. I have to read it twice to understand.. Therefore it was mainly due to words we used like 'misery'. In addition to that, positive feedback were also given, such as: "It shows me that it knows that it's not my favorite song but also tries to minimize misery.", or "At least it explains the reasoning."

It can be interpreted as people are prefer to receive more friendly and light explanations rather than explanations with complicated words to describe the aggregation strategy behind the sequence generation.

In addition to that, explanation 2 in the same scenario (scenario 2) performed slightly better ($m=2.81$, $SD=1.05$), as this more complete explanation contains ratings which helped users understand the explanation better.

Users' Comments

Additional comments related to individual participants are as follows: "It depends on my personality and mood. Example: if I am in my car with friends in summer and there is sunshine I could be happily let others favorite songs play even if I hate that. But if it's winter and I'm sad, I can't accept it easily.". Other comments about complete explanations vary between participants e.g., some have comments like "Good to know about the ratings." but on the other hand for the same explanations others have comments like "My friends rating is not so interesting, it's sort of privacy violation." or "I would not be comfortable with the system giving out my rating". Users' comments highlighted the need for privacy when revealing their personal information, like their preferences/ ratings in the group.

In this experiment, we investigated how different explanation styles work for different settings (depending on the preferences composition in the group). Based on the results, we saw the importance of adapting the text of the explanation as well as more general guidelines (i.e., brevity and clarity). In this chapter, we also evaluated aggregation strategies and their explanations. The following section (see Section 3.4) summarizes our findings of these three experiments in this chapter.

3.4 Chapter Conclusions

In a group recommendation/decision context, there is information that we can present to people in a group to justify why they should follow certain recommendations. In this Chapter, we covered three experiments looking at the content of explanations. The aggregation strategies experiment evaluated different aggregation strategies used to generate group recommendations. The results suggested that we probably do not need to consider which strategy we explain as long as we explain the used strategy. So to understand whether the explanation is necessary or not in the aggregation strategies versus explanations experiment, we studied the effect of the aggregation strategy separately from the effect of the group explanation. Our findings suggested that the need for explanations may depend on the complexity of the recommendation scenario (preference/item rating agreement or disagreement in the group). In a follow-up experiment (the formulating group explanations experiment) we proposed and evaluated different explanations styles for the aggregation strategies for more complex scenarios (depending on preferences composition in the group).

Experiment 1: aggregation strategies. It is not clear which information, in general, we should convey to the group to support them. As for single-user recommendation explanations, explanations for groups can be designed based on the underlying recommendation algorithm. There are underlying methods for deciding what we can recommend to groups; the choice of the used algorithm also may influence the types of explanations we can generate. We started with evaluating different aggregation strategies. We presented a user evaluation of four different explainable aggregation strategies, namely Least+, Fair+, Average, and Dictatorship, in two scenarios (groups with different preferences versus groups with similar preferences) in the *tourism* domain.

We found a significant difference between algorithms in terms of the combined variables (perceived individual and group satisfaction, fairness, and acceptance). Further analysis showed a difference between the Dictatorship strategy and the other strategies. The Dictatorship strategy scored lower compared to the other three strategies in terms of both user-perceived individual and group satisfaction. But there was no other differences between strategies. User comments strengthen this finding that our participants were sensitive to the dictator-based strategy, which (comparatively, negatively) affected their satisfaction on their behalf and behalf of the group. This suggests no matter which strategy we use, as long as it is not a dictator strategy that only considers the person in the authority's preferences. In other words, it does not seem to matter in terms of user satisfaction which strategy we apply. *So we probably do not need to consider which strategy we explain as long as we explain the used strategy.* So one way to understand whether the explanation is necessary or not in the first place is to separate the output of the aggregation strategy from the output of the explanation, which was investigated in the second experiment.

Experiment 2: aggregation strategies versus explanations. In the second study, we evaluated the effect of the aggregation strategy separately from the effect of natural language explanation. The results do not show a significant benefit of using explanations, at least in simple scenarios (i.e., when there is a few recommended items, only three items in our case) like the one we used. This needs a further investigation to see in what scenarios

explanations are more needed. Overall, our findings suggest the need for explanations may depend on the complexity of the recommendation scenario (preference/item rating agreement or disagreement in the group) and domain. We also discussed some of the challenges and decision points required to benchmark future studies of group explanations. In particular, we highlighted the importance of clarifying and motivating the recruitment process and properly choosing the experimental design, specifying how each condition is assigned to each participant. Furthermore, we highlighted how the choice of the scenario to present for the evaluation could influence the results. Therefore, the results should always be discussed in relation to the scenario, i.e., preferences/ratings agreement or disagreement in the group.

3

Experiment 3: formulating group explanations. In the third experiment, we then suggested different explanations styles for the aggregation strategies for more complex scenarios (depending on preferences composition in the group). For example, *reassuring* style for when the group agrees on the recommended item or *repairing* style for when the group disagrees on the recommended item. Participants preferred short (only represents the ratings of the members that the item is recommended based on those), encouraging explanations rather than detailed (represents the ratings of everyone in the group), and negative explanations. However, a more detailed explanation seems more beneficial for users when maximal misery (not getting their liked item)/group disagreement was expected. However, based on qualitative results, privacy (concerns about the possible consequences of disclosing a specific piece of information in a group explanation) stands out as an influential factor when generating detailed group explanations. This remains for the next chapter to investigate this important aspect that which information people are more sensitive to disclosing in a group explanation.

Wrap up

- No matter which aggregation strategy we use to develop group recommendations, as long as it is not a dictator strategy. As this strategy only considers the preferences of the person in the authority. This appears to be the case for settings with a lot of similarity in the users' preferences and a low to medium involvement domain. So in these situations, we probably do not need to consider which strategy we explain as long as we explain the used strategy.
- There does not appear to be a benefit of adding textual explanations to show the results of an aggregation strategy. This seems to be the case in a more artificial/basic scenario with a low number of items (three in our case) and a small group size (three in our case). However the limited benefit of such explanations cannot be generalized for real-world scenarios.
- There appears to be a benefit of modifying the group explanation style (i.e., repairing vs. reassuring) to the variation in user preferences (i.e., when there is group disagreement or agreement on the recommended item). For example, when group members have different preferences/choices, using repairing style explanations that

focus on why the item is recommended and include more details seems more beneficial. However, users' comments suggest that this leads to privacy issues concerning which types of information can be mentioned in an explanation (i.e., some users were sensitive to disclosing their preferences/item ratings in a group explanation).

Based on these findings, we do not continue to include aggregation strategies description in explanations for groups and focus more on the trade-offs between group members' need for privacy and transparency in explanations of group recommendations in the remaining of the thesis.

4

Privacy in Group Explanations

4

My thesis started with the aim of studying what makes good explanations for group recommendations. Imagine a group of people trying to make a joint decision, for example, a group of colleagues on a business trip deciding where to visit. Various strategies can be used to predict an item suitable for the group. Reaching a consensus in the group can be challenging, especially when group members have different preferences. In this situation, explanations can help people understand how certain recommendations are generated. Explanations can be regarded as additional information (i.e., in this thesis this information is textual only) that accompanies the recommendations and serves various goals, such as increasing satisfaction (the ease of usability or enjoyment of the used recommender system) [124]. However, among the many types of information to include and many ways to formulate an explanation, it is unknown which information should be shown in the explanation. As a first step, in Chapter 3, we started investigating what information should be conveyed when generating group explanations. User comments highlighted that people might be sensitive to disclosing some of their information, e.g., their preferences. This leads us to the second research question, which investigates the factors that we should model in the group to consider group members' privacy risk of information disclosure (RQ2). We conducted three experiments discussed in this chapter to answer this research question. The first experiment (the privacy preferences experiment) investigates which information people would like to disclose or not disclose in explanations for group recommendations (see Section 4.1). We found that certain types of information are more sensitive than others, but we also see reasons to believe this varies between individuals. We wanted to understand what contributes to these individual differences. To understand this better, in the second experiment (the privacy factors experiment, see Section 4.2), we studied some factors identified in the literature and the previous experiment (the privacy preferences experiment) that influence individual privacy risk, i.e., group members' personality, the type of relationship they have in the group, and preference scenario (whether their preferences are aligned or not aligned with the preferences of the majority in the group). We studied these factors when disclosing several types of personal information in a single group explanation (i.e., location, drug/alcohol, emotion, personal details, and personally identifiable information). These information categories are a subset of nine personal information categories that are relevant to the domain of tourism. The results help to see whether these factors

generally affect individual privacy risk or not. As we found significant effects of these factors on users' privacy risk perception, in the next experiment (the information disclosure experiment), we studied how these factors affect group members' actual information disclosure (see Section 4.3). Although we have expected to see the opposite effect of factors (compared to the privacy factors experiment's result) on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure), neither the personality traits nor the preference scenario affected people's actual disclosure. One explanation for why we did not obtain the expected results in the information disclosure experiment could have been that our task design nudged participants into a "convincing mindset". This could have sheltered the effects of previously identified factors on actual disclosure. Further investigation on the effect of task design on actual information disclosure in the group remains for the next chapter.

4

This chapter is based on a workshop paper and two conference papers:

- Shabnam Najafian, Oana Inel, and Nava Tintarev. *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, pages 85–86, 2020.
- Shabnam Najafian, Amra Delic, Marko Tkalcić, and Nava Tintarev. *Factors influencing privacy concern for explanations of group recommendation*. In Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization, pages 14–23, 2021.
- Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcić, Jie Yang, and Nava Tintarev. *Exploring user concerns about disclosing location and emotion information in group recommendations*. In Proceedings of the 32nd ACM Conference on Hypertext and Social Media, pages 155–164, 2021.

4.1 Experiment 1: Privacy Preferences

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 make it easier for users to accept items they might not like for the benefit of the group [127]. Many studies have demonstrated the benefits of adding explanations to automated recommendations (i.e., [42]). The majority of research focused on single-user scenarios. However, when explaining recommendations to a group of users, an additional aspect, *privacy*, becomes relevant as well. This aspect requires a trade-off between a) generating effective explanations to group members and b) keeping each group member comfortable by not disclosing private information, e.g., their preferences, to other group members.

This raises the question of *which information an explanation should disclose when displayed to the whole group*. To answer this questions, in this section, we (dynamically) generate natural language explanations for group recommendations (see Section 4.1). The study design allowed us to compare users' privacy preferences for different low consensus scenarios, where either the active user or their acquaintances did not get their preferred item, with a high consensus scenario (where both the active user and acquaintances get their preferred item). The generated explanations are evaluated in a within-subjects user study (n=200) where users are able to specify their preferred explanation settings (Section 4.1.2). We describe the results and discuss the limitations of our study in Section 4.1.3. Finally, we discuss some practical implications and conclude with plans for future work in Section 4.4.

This study provides the following contributions:

- We designed a system to (dynamically) generate and adapt natural language explanations in the context of group recommendations;
- In a user study (n=200), we studied user's privacy preferences regarding the generated explanations in two low consensus scenarios; **a)** where either the active user or **b)** their acquaintances did not get their preferred item; and a high consensus scenario: **c)** where all group members got their preferred item);

We found that people were generally willing to disclose a lot of information. However, we found that people prefer more private explanations (to hide more personal information) in both low consensus scenarios than in a high consensus scenario. One surprising result was that people avoided disclosing the combination of name and personality in all the scenarios. In line with previous work, [62], our findings suggest that there may be some individual differences in the levels of participants' concern perception about being singled out for having different preferences. Whether we can predict which factors (e.g., personality, relationship, etc.) contribute to these differences led our following study to investigate this further (see Section 4.2).

The contribution of this study is published as a short paper in Proceedings of the 25th International Conference on Intelligent User Interfaces Companion [86].

4.1.1 Preliminary Definitions: Generating Natural Language Explanations

In this section, 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, and personality. For instance, if users decide to, for example, hide their names but show their 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 [110]:

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 on how group members' preferences and personalities played a role in generating the recommended item. 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 [131].¹ *"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. [131] and the part for personality from Quijano-Sanchez et al. [108]. Below we use a working example for the scenario where the active user did not get their preference, but their acquaintances did.




Figure 4.1: Screenshot of the system. Users can adjust the generated explanation using three different privacy controls: name, rating, and 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 item, but her acquaintances do.

1. **Name:** we picked parts of the template from [131] 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

¹This 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.

(e.g., *"This decision does not support the preferences of Ana .."*). Note: The positive and negative parts depend on the scenario.

2. **Rating:** In a previous pilot study [107], 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 item 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. [108], we only show the personality of assertive members who have strong opinions and are difficult to convince. The member(s) with a strong opinion is always assumed to be the same users who got their preferred item. 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* [23]. 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 SimpleNLG² 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 4.1.

4.1.2 Experimental Design

We wanted to understand how much information the recommender system should expose to the group in different low consensus scenarios compared to a high consensus scenario.

To answer this question, we conducted a within-subjects online experiment with 200 participants recruited from the Amazon Mechanical Turk (mTurk) crowdsourcing platform.³

The participants were asked to imagine they are carpooling with two acquaintances and listening to music during a road trip. They were told that the same explanation is shown to all group members. Since people might choose different privacy options in different scenarios, in a within-subjects design, each participant completes the questionnaire for three scenarios (c.f., Section 4.1.2). E.g., *Imagine that you are carpooling with*

²SimpleNLG (v. 4.4.8) is a "realization engine" built by Albert Gatt and Ehud Reiter [34].

³<https://www.mturk.com>.

two acquaintances and you are being recommended a song that you don't like, but the two acquaintances like it. ('Unhappy User').

To understand how much information the recommender system should expose to the group we asked users to adjust the explanations with the information they feel comfortable to share with their group members.

They were able to control three privacy-related option/parameters (namely, whether to show names, ratings, and personality) to adjust the explanation. Figure 4.1 demonstrates these three options when they are all selected to be on. For example, when the name button is on, the explanation includes the group members' names. If the name button is off, it anonymizes names with 'some group members' and 'all group members'.

Once participants were satisfied with the generated explanation, they were asked to evaluate the explanation (for transparency, privacy threat, local privacy, and satisfaction).

4

Independent Variables

In this experiment, we investigate the use of explanations with different privacy control options (privacy adaptations) for group recommendations.

Types of privacy information. Our aim is to give users control over which information they share with the group, from the following options:

- *Name*: show or hide the names of the group members whose the system recommendation supported their preference or not supported their preference. With the name button *on* the explanation includes group member names. Otherwise, if they select this button *off* names will be replaced with anonymous names like some users or all users;
- *Rating*: show or hide all group members' ratings regarding the recommended item. If they select the rating button *on* the explanation will include *who did not rate this song highly* or *who rated this song highly*. Otherwise, if they select this button *off* these sentences will be excluded from the explanation;
- *Personality*: show or hide all group members' personalities. If they select the personality button *on*, the explanation will include the sentence *Besides, we have detected that users A and B really want to listen to this song and won't be talked out of it easily*. Otherwise, if they select this button *off* the sentence will be excluded from the explanation;

Scenarios. People might choose different privacy options in different scenarios. For instance, when they do not get their preferred item, they might need a more transparent explanation to comprehend the system's decision-making. However, when the user gets her preferred item but the other two group members do not get their preferred item, a user might not feel comfortable when the system discloses her name in the explanation.

The literature suggests that relationship proximity influences the importance of coherence in a group [135]. In all the scenarios, the participant is told that they are with a group of acquaintances. We chose this to represent a loosely coupled group, as an earlier pilot study found that people in this type of group might care more about their privacy.

The literature also suggests that the level of privacy needed will depend on the scenario (context) [62]. We are therefore interested in understanding which privacy control options are more likely to be picked for the following predefined scenarios:

1. The user **does not get** her preferred item, however, her acquaintances **get** their preferred item. ('unhappy user')
2. The user **gets** her preferred item, however, her acquaintances **do not get** their preferred item. ('unhappy acquaintances')
3. The user **gets** her preferred item, and, her acquaintances **get** their preferred item. (baseline)

Dependent Variables

The selection made by each user was dummy coded into two dichotomous values for each information type: **name, rating, personality**. For example, a selection of show names was coded "1", and the hide name as a "0".

Additional Variables

To support qualitative analysis, we also evaluated the generated explanations in terms of the following criteria:

- **Transparency:** how much the explanation helped users to understand why the system recommended the item it did;
- **Privacy threat:** how unsafe users felt regarding disclosing all group members' personal information to the other group members, through the generated explanation;
- **Local privacy:** how users felt regarding disclosing other group members' personal information but not her personal information through the generated explanation;
- **Satisfaction:** how much the explanation helped increase the group members' satisfaction with regards to the recommended item.

(with a 5-point Likert scale):

- Q1: I understand why the system recommended the song it did [16] (transparency).
- Q2: I understand what the system bases its recommendations on [16] (transparency).
- Q3: The system disclosed information about me that I consider private [62] (privacy-threat).
- Q4: I felt tricked into disclosing more information about myself to the group than I wanted [62] (privacy-threat).
- Q5: The system shows more about me than I am comfortable with [62] (privacy-threat).
- Q6: I am OK with the system disclosing the information of others but not mine (local privacy).

- Q7: I am OK with the system disclosing my information but not the others (local privacy).
- Q8: I would not like this explanation to be shown to other people in the group (local privacy), (reversed scale).
- Q9: The explanation helps to increase the satisfaction of group members with regard to the group recommendation [131] (satisfaction).

Hypotheses

- **H1)** People need more transparent explanations when the user does not get her preferred item ('unhappy user') than in a high consensus scenario (baseline).
- **H2)** People need less transparent explanations when the acquaintances do not get their preferred item ('unhappy acquaintances') than in a high consensus scenario (baseline).

4

Statistical Analyses

We wanted to determine if there are nonrandom associations (with regard to the selected options) between two categorical variables (in our case, two scenarios). We applied Chi-square for within-subjects (McNemar-Bowker test) to test our hypotheses. Bonferroni correction was applied when multiple tests were conducted. We calculated the minimum number of required sample size (i.e., the number of study participants needed) based on the G*Power analysis for Chi-square for within-subjects (McNemar's test) [26]. The analysis showed that we needed at least 200 participants.

4.1.3 Results

In this section, we describe the results of users' privacy choices for the explanation of recommendations in the context of groups for three scenarios. For brevity, we will henceforth call these the '*unhappy user*' scenario (user does not get their preferred item, but their acquaintances do), the '*unhappy acquaintances*' scenario (acquaintances do not get their preferred item, but the active user does), and the *baseline* (everyone gets their preferred item) respectively.

Participants

We selected master workers from mTurk, with an approval rate higher than 98%, and only from native English-speaking countries (United States, United Kingdom, Canada, and Australia). We paid each worker 2 USD. Besides, we used 'honeypot' questions [72] to filter untrustworthy workers. For this purpose, we randomly inserted a set of trapping questions (whose true answer is already known) (5 workers excluded). In the next step, we also excluded users whose answers to the free text question regarding their motivation of the selected options did not make sense (10 workers).

This resulted in 200 participants, 42% female, and 58% male. The highest level of education that they hold was 26% a high school diploma or equivalent degree, 55% a bachelor's, and 19% a master's degree or higher. Among those, 68% use music applications (such as Spotify) at least every day and 24% at least every week, and only 9% use it less than once a month or never.

Table 4.1: Selection of the three choices is summarized separately. Values represent frequencies of the chosen privacy option to show: name, rating, and personality in each scenario and overall. Percentages are given in parentheses, and the maximum and minimum frequencies are in bold.

CONDITION	PRIVACY OPTIONS		
	Name	Rating	Personality
Baseline	136 (68%)	149 (74%)	113 (56%)
Unhappy User	108 (54%)	140 (70%)	105 (52%)
Unhappy Acquaintances	99 (49%)	126 (63%)	106 (53%)
Overall	343 (57%)	415 (69%)	324 (54%)

Table 4.2: Selection of the three choices is summarized by eight patterns. Values represent frequencies of the chosen privacy control options: (no) name (nm), (no) rating (rtg), and (no) personality (prs) in each scenario and overall. Percentages are given in parentheses, and the maximum and minimum frequencies are highlighted in bold.

Privacy Control Options	Unhappy User	Baseline	Unhappy Acquaintances	Overall
1) no nm/ no rtg/ no prs	23 (11%)	13 (6%)	26 (13%)	62 (10%)
2) no nm/ no rtg/ prs	12 (6%)	10 (5%)	23 (11%)	45 (7%)
3) no nm/ rtg/ no prs	32 (16%)	28 (14%)	32 (16%)	92 (15%)
4) no nm/ rtg/ prs	25 (12%)	13 (6%)	20 (10%)	58 (9%)
5) nm/ no rtg/ no prs	20 (10%)	19 (9%)	18 (9%)	57 (9%)
6) nm/ no rtg/ prs	5 (2%)	9 (4%)	7 (3%)	21 (3%)
7) nm/ rtg/ no prs	20 (10%)	27 (13%)	18 (9%)	65 (11%)
8) nm/ rtg/ prs	63 (31%)	81 (40%)	56 (28%)	200 (33%)

Overall Preferences

To gain insight into users' privacy preferences, we first look at their overall choices. We want to understand which options users are more likely to disclose across scenarios. Table 4.1 demonstrates the results for each scenario and overall, for each privacy option. Participants overall preferred to show their personality less often than the other two privacy options (54%). After disclosing personality, disclosing their name was the least frequent option people chose (57%). The results show overall, disclosing the rating was the most selected option among the provided privacy options (69%).

To better understand the way people combined privacy options, we analyzed these in Table 4.2. In this table, each row represents a unique pattern of showing or hiding the three different types of information, e.g., **no nm/ rtg/ no prs** reflects the pattern where users show rating while hiding name and personality. The last column ("Overall") shows the aggregated choices for all three scenarios. The most frequently selected pattern was the one where the participants disclosed *all* their personal information, including name, rating, and personality (200 choices out of 600, 33%). We further analyze the patterns for the individual types of private information.

Name. As can be seen in Table 4.2 column "Overall", hide name options have been selected more than their equivalent combination with show name option (patterns 1 vs 5; 2 vs 6; 3 vs 7). The most selected pattern is disclosing *all* the information: name, rating, and personality (pattern 8). The least selected pattern is the combination of (show) name and (show) personality together (pattern 6). This will be discussed further in Section 4.1.4.

Rating. As can be seen in Table 4.2, column "Overall", all (show) rating options are selected more than their equivalent combination with hide rating option (patterns 1 vs 3; 2 vs 4; 5 vs 7; and 6 vs 8).

Personality. As can be seen in Table 4.2, column "Overall", most hide personality options were selected more than their equivalent combination with show personality option (patterns 1 vs 2; 3 vs 4; 5 vs 6). The combination of name and personality was chosen the least as mentioned above. This surprising result will be discussed further in Section 4.1.4.

H1: 'Unhappy Users' will Want more Transparent Explanations

According to H1, we hypothesize that for the 'unhappy user' scenario (s)he will want more transparent explanations compared to the baseline scenario where all group members get their preferred item. So we expected users to select more (show) names, (show) ratings, and (show) personality options compared to the baseline. However, in Table 4.1 we see that the frequency of showing private information is *lower* for the 'unhappy user' scenario than the baseline for all three privacy options.

While we found a statistically significant difference between the two scenarios ($p < 0.05$, two-sided McNemar-Bowker test), the results are in the opposite direction predicted by H1, that is, **people seem to prefer less transparent explanations in the 'unhappy user' scenario, than in the baseline.** This is confirmed by user comments which we will discuss further in Section 4.1.4.

To better understand the way people combined privacy options, we analyzed these in Table 4.2. We compare the first ("Unhappy user") and second ("Baseline") columns of Table 4.2 where we report the frequency for each of the privacy patterns.

We observe that only 63 users out of 200 (31%) selected completely transparent explanations (pattern 8) in the 'unhappy user' scenario, while 81 users out of 200 (40%) selected completely transparent explanations in the baseline.

We can also see that 23 users out of 200 (11%) selected completely private (pattern 1) explanations in the 'unhappy user' scenario, while only 13 users out of 200 (6%) selected completely private explanations in the baseline.

We further analyze the patterns for the individual types of private information.

Name. We see that participants chose to *hide name* more often in the ‘unhappy user’ scenario than in the baseline (patterns 1, 2, 3, and 4). Analogously, in the ‘unhappy user’ scenario, people also selected the *show name* option less compared to the baseline (patterns 6, 7, and 8).

Personality. In Table 4.2, comparing columns “unhappy user” vs “the baseline” shows, people selected the hide personality option more in the ‘unhappy user’ scenario compared to the baseline in patterns 1, 3, and 5 but not in pattern 7. We see that the show personality option in combination with the show name has been selected less frequently than in the baseline scenario for patterns 6 and 8, and show personality in combination with hide name has been selected more frequently in patterns 2 and 4. This suggests an issue with showing personality and name together, which we will address in Section 4.1.4.

Rating. In Table 4.2, comparing columns “Unhappy user” vs “the baseline” shows that people selected the hide rating option more in the ‘unhappy user’ scenario compared to the baseline scenario (patterns 1, 2, 5). When show name was selected, participants selected hide rating more often compared to the baseline (patterns 6, 7, and 8). Overall, it looks like the effects of the show name are so strong that it affects the selection of show rating.

H2: With ‘Unhappy Acquaintances’ People Want Less Transparent Explanations

According to H2, we hypothesize that for the ‘unhappy acquaintances’ scenario (s)he will want less transparent explanations compared to the baseline scenario to, for example, avoid being disliked in the group. We expected users to select fewer show names, show ratings, and show personality options compared to the baseline.

In Table 4.1 we can compare the frequency of the selected privacy options in the ‘unhappy acquaintances’ scenario and the baseline. We see that all three privacy options (show name, show rating, and show personality) were selected less frequently than for the baseline. This also reflects a statistically significant difference between the two scenarios ($p < 0.05$, two-sided McNemar-Bowker test). These results appear to confirm H2, that is, **people seem to prefer less transparent explanations in the ‘unhappy acquaintances’ scenario than in the baseline.**

To better understand the way people combined privacy options, we analyzed these in more detail in Table 4.2. We compare the third (“Unhappy acquaintances”) and second (“the baseline”) columns, where we report the frequency of the privacy patterns users selected to be disclosed for these two scenarios.

Fewer users (56 users out of 200, 28%) selected completely transparent explanations (pattern 8) in the ‘unhappy acquaintances’ scenario than in the baseline (81 users, 40%). We can also see that more (26 users, 13%) selected the completely private (pattern 1) explanations in the ‘unhappy acquaintances’ scenario, compared to the baseline (13 users, 6%).

Name. We see that people selected the hide name option (patterns 1, 2, 3, and 4) more in the ‘unhappy acquaintances’ scenario compared to the baseline scenario. Similarly, they selected the (show) name option less frequently compared to the baseline.

Personality. The combination of name and personality appear to interact, showing one more correlate with hiding the other one more. Either hide or show personality with hide name (patterns 1, 2, 3, and 4) have been selected more in the ‘unhappy acquaintances’ scenario compared to the baseline. The combination of either hide or show personality with show name (patterns 5, 6, 7, and 8) has been selected less frequently in the ‘unhappy acquaintances’ scenario compared to the baseline. As for the ‘unhappy user’ scenario, this suggests an issue with showing personality and name together. We address this interaction in Section 4.1.4.

4

Rating. As in the ‘unhappy user’ scenario, we can see the effect of the name is strong, which overshadows the effect of the rating. The combination of show name and hide rating, as well as the combination of hide name and show rating, are selected more in the ‘unhappy acquaintances’ scenario compared to the baseline.

Qualitative Feedback (comments)

Participants were asked to give free text comments to motivate their choices. In this subsection, we report on the participants’ motivation regarding their selected privacy options.

The ‘baseline’ scenario. In this scenario, every group member gets their preferred item. Although some users mentioned it was OK to disclose more personal information, others still felt insecure about disclosing personal information. For example, they mentioned, *“I don’t like apps tracking my name or personality in any way”* or *“Names being shown felt less like a privacy violation in this scenario. Still, I’m a very private person and I’d just as soon that option wasn’t available”*.

The ‘unhappy user’ scenario. In the ‘unhappy user’ scenario, the active user does not get his/her preferred item, but his/her acquaintances get their preferred item. In this case, participants who chose more private explanations mentioned they did that to avoid conflict with other group members and to seem they liked the same item as other group members. They did not want to be singled out as not liking the item.

For example, *“Telling who exactly likes/rated highly feels like a violation of privacy. Maybe I would like for my friends to think that I like the same song that they like. I don’t want to be singled out as “that guy.”*

However, other users liked to see who liked the item to decide whether they were in the majority or minority in terms of their preferences for the item. If the majority dislike the item, then they skip the item to something everybody likes. For example, *“I wanted to see who liked the song in order to make a decision about whether I was in the majority or minority for thoughts about the song. If the majority of the car liked the song, then it should be played. If the majority dislike the song, it should be okay to skip to something everybody likes”*.

Some participants stated that they hid personality information because it described their group members negatively. Others hid personality information because it is not clear how the system would have come to that conclusion. For example, *“I hid personalities because it described my friends negatively, and I don’t understand how it would have come to that conclusion. It’s an opinion and not something my friends chose to share. Keeping names and ratings showing is fine to me because it’s objectively showing how everyone feels about the song based exactly on what each person rated and not opinions of the system”*.

Those who selected more transparent explanations mentioned that the information that the system has is not sensitive to them and that they do not care if their music preferences are disclosed. For example, *“This situation didn’t seem to violate on standards or norms I have on privacy. If I’m carpooling with people, I assume they can see that I am male, 38 years old, white, etc. I don’t see song preferences and ratings as highly sensitive information”*.

The ‘unhappy acquaintances’ scenario. In the ‘unhappy acquaintances’ scenario, the active user gets his/her preferred item, but his/her acquaintances do not get their preferred item.

One of the motivations users had for selecting more private explanations was they would rather that the whole group thinks that the majority of people want to listen to it. For example, *“I’d rather that no one knows why the song is being played. I rather they all think that the majority of people want to listen to it”*.

The users who elected to only show ratings stated that they felt that hiding the names helped to not single themselves out. For example, *“I felt that hiding the names helped to not single out a user. Me in this case”*.

Other users who picked more transparent explanations were OK with the system justifying the selection due to their preferences:

For example, *“I wanted to include that even though no one enjoyed it, I enjoyed this music so that is a very reasonable explanation why I had the song play”*. This suggests that there are some differences in the levels of participants’ concern about being singled out as a minority.

Additional Variables

We additionally evaluated the generated explanation based on users’ perceived transparency, privacy threat, local privacy, and satisfaction regarding the recommender system. These are meant as descriptive measures. Table 4.3 shows the average and standard deviation of participants’ responses to these four measures. We expect a high value for transparency and satisfaction and a low value for privacy threats. We also expect higher values for Q6

Transparency. As can be seen in Table 4.3, the mean and standard deviation of user responses for perceived **transparency** for both transparency questions (Section 4.1.2: Q1, Q2) were quite similar in all the scenarios. The results show overall perceived transparency was high in all three scenarios. Comparing between scenarios, we see the participants’ perceived transparency for the baseline is comparable to the other two scenarios (unhappy user, unhappy acquaintances).

Table 4.3: Mean and Standard Deviation (SD) of participants' responses to the following measures: perceived transparency, privacy threat, local privacy, and satisfaction per scenario. For the responses, a 5-point Likert scale was used with 1= "strongly disagree" to 5 = "strongly agree".

Criteria	Question	Unhappy User		Unhappy Acquaintances		Baseline	
		Mean	SD	Mean	SD	Mean	SD
Transparency	Q1	3.9	0.9	3.8	1	4.1	0.9
	Q2	3.9	0.9	3.9	0.8	4.1	0.7
Privacy threat	Q3	2.4	1.2	2.4	1.2	2.3	1.2
	Q4	2.5	1.4	2.4	1.3	2.3	1.3
	Q5	2.5	1.3	2.5	1.3	2.4	1.3
Local privacy	Q6	3	1.2	2.9	1.1	2.9	1.1
	Q7	2.9	1.3	3	1.2	2.8	1.1
	Q8	2.5	1.2	2.7	1.2	2.5	1.4
Satisfaction	Q9	3.9	0.8	3.8	0.9	4.1	0.8

Privacy threat. Users' responses to the three questions of perceiving **privacy threat** (Section 4.1.2: Q3, Q4, Q5) were quite similar as well. This shows participants perceived a relatively low risk of disclosing their private information. This applies to all three scenarios. Comparing between scenarios, we see users' perceived privacy threat for the baseline is slightly lower than the other two scenarios (unhappy user, unhappy acquaintances), which means that users perceived a lower risk of disclosing information in the baseline scenario.

Local privacy. Participants' responses to the **local privacy** questions in the 'unhappy user' scenario and the baseline scenario show a preference to protect their own information (Q6) but not the privacy of others (Q7).

In contrast, in the 'unhappy acquaintances' scenario, participants' preferred to show their own information but not other members' information (Q7) more than showing their own information but not information for other members (Q6).

In all three scenarios, participants did not feel completely comfortable sharing the presented explanation with the group (Q8).

Satisfaction. Participants' responses show that they, on average, and in all scenarios, perceived high **satisfaction** after getting an explanation for the recommended item.

4.1.4 Discussion

Overall willingness to disclose personal information. As can be seen in the results, people were generally willing to disclose a lot of information. This preference was higher in the high consensus scenario (the baseline) rather than two low consensus scenarios. Motivated by user comments, even in the ‘unhappy acquaintances’ scenario in which we expected that people try to select the less transparent explanation, 28% still selected the full transparent explanation to, for example, persuade other group members to listen to their preferred song. For example, *“I chose full transparency explanation as it tries to persuade the others why to let me listen to the song”* or *“I liked that it gave the most information about why it was selecting a song. The fact that I liked it and would not be easily talked out of it would help to satisfy the other people”*.

Willingness to hide personal information in low consensus scenarios. We found that in both low consensus scenarios (where either the active user or their acquaintances did not get their preferred item), users chose to hide all personal information more than in a high consensus scenario (where all group members get their preferred item). However, comparing the two low consensus scenarios, we can see that the ‘unhappy acquaintances’ users tend to hide personal information more than in the ‘unhappy user’ scenario.

In contrast to our hypothesis, ‘unhappy users’ still preferred to hide their dislike of the recommended item compared to the baseline in order to e.g., avoid conflict in the group. This is motivated by user comments; they did not want to be singled out in a conflict situation. For example, *“I just don’t want to have any conflict with anyone else. I’d rather just listen to what the majority of people want to, even if I don’t like it myself”*.

For the ‘unhappy acquaintances’ scenario, we saw that users, in line with our hypothesis, preferred to hide the fact that the item was recommended because they liked that. For example, *“I’d rather that no one knows why the song is being played. I rather they all think that the majority of people want to listen to it”*.

As can be seen in Table 4.1, in both low consensus scenarios (‘unhappy user’, ‘unhappy acquaintances’), users wanted to hide personal information more than the baseline.

Avoid disclosing name and personality together. Overall, people did not want to disclose names and personalities, also effect was strongest for the unhappy user. We also found that in the scenario that benefited the active user, they preferred to show that there is a strong opinion (assertive personality) as long as the system does not disclose their own identity. For example, *“I guess it’s nice to feel justified in hearing a song you like when others don’t if it’s anonymous. You don’t have to be embarrassed, but you still get your preference attended to. The personality bit is cuter than I would have thought”* or *“I like the personality option because it can help one person that likes a song not feel outnumbered when the rest don’t”*.

These results have implications in particular for shared explanation interfaces. For example, when the minority gets their preferred item, we should be cautious about disclosing names together with strong opinions.

Additional variables. After evaluating the user-adapted explanations, the high values for user perceived transparency and satisfaction and low values for user perceived pri-

vacuity threat, and local privacy shows that the privacy options that we provided were good enough, and participants were able to generate an explanation that met their privacy expectations.

Comparing between scenarios, we see users' perceived higher transparency and lower risk of disclosing information (privacy threat) for the baseline compared to the other two scenarios (unhappy user, unhappy acquaintances), which is as expected.

Limitations

Our experiment was designed as a controlled experiment, which has allowed us to study the trade-off between transparency and privacy of explanations for group recommendations. However, it has constrained us in studying other relevant factors. We discuss these limitations below.

4

Experimental setup. In this study, we were only able to measure the preferences of the active user. However, this study considers their privacy preferences in three scenarios, which allows us to make interesting inferences about different types of low consensus situations compared to a high consensus scenario.

Group size. For the purpose of this experiment, we restricted the group size to exactly 3 participants, *i.e.*, an active user, and two acquaintances. We would expect privacy considerations to be more extreme in larger groups, as well as more dependent on group dynamics. There is also limited consensus on how to measure satisfaction within a group (E.g., Do we measure the average satisfaction? Do we only want to avoid very low ratings? etc), which should be developed in tandem with such an experimental design.

Hypothetical recommendations. This study focused solely on the different explanations evaluation, and we have not fully explored the role of these explanations in a live recommendation setting in groups. This means a lower ecological validity of the study, but it also allowed us to control for many external factors such as familiarity with specific songs, visual appeal of albums, etc.

Domain specific results. Like many previous evaluations of an adaptive system, we are evaluating an adaptation in one particular domain (music) [126]. This is a low-risk domain, and it is more difficult to say if they would have chosen to disclose less information in high-risk domains. This point was also mentioned in the user comments, e.g., "*The information is not at all sensitive. Why would I care if people know what music I like?*" or "*... I don't see song preferences and ratings as highly sensitive information.*"

Negative formulation of personality. We made some simplifying assumptions about the formulation of the natural language explanations. Most notably, the part focused on personality focused on (a) group member(s) with strong opinions. This is based on previous work which recommended basing explanations on strong opinions, which found that saying some group members had weak opinions was not tactful [108]. However, user comments in our study suggest that many users found this part of the explanation helpful,

e.g., “... *The fact that I liked it and would not be easily talked out of it would help to satisfy the other people*”.

Personality is a derived feature. Unlike rating and names, the personality of participants is somehow learned or calculated. This means that we cannot guarantee its accuracy, and this, in turn, affects user acceptance of this part of the explanation. This may be part of the reason that the frequency for the personality information was lower than for name and rating. This was also mentioned in user comments: “... *And I felt that personalities was an assumption, at best.*” or “... *Furthermore, with the personality rating, I don't think that there is a lot of accuracy to that, and I'd probably scrap it.*” However, we also see that many participants also chose to disclose personalities, even for themselves. The personality information was hidden more often when combined with the name.

This experiment investigated which information people would like to disclose in explanations for group recommendations. We found that certain types of information are more sensitive than others, but we also see reasons to believe this varies between individuals. We wanted to understand what contributes to these individual differences. For example, the effect of the preference scenario (i.e., whether group members' preferences are aligned or not aligned with the preferences of the majority in the group) on perceived privacy risk. User comments suggest that there are some differences in the levels of participants' concern about being singled out as a minority. To understand this better, in the next section (see Section 4.2), we investigate the relationship between factors identified in the literature and highlighted in this experiment that affects individuals' perceived privacy risk, namely: people's personality, their relationship type in the group, and preference scenario when explaining group recommendations.

4.2 Experiment 2: Factors that Affect Privacy

This thesis started with the aim of studying what makes good explanations for group recommendations. In Chapter 3, we started by determining what information needed to be conveyed to people in a group explanation for the recommended items. User comments highlighted the need for protecting certain types of information when presenting an explanation to the group, i.e., group members' ratings of items. Besides, findings from previous experiment suggest that there may be some individual differences in the levels of participants' privacy risk of disclosing their information (the privacy preferences experiment, see Section 4.1). To understand what contributes to these individual differences, in this experiment, we investigate the relationship between factors identified in the literature of individual privacy risk and also on the basis of the previous experiment (the privacy preferences experiment). The first factor we investigate is users' *personality*, modeled using the Five Factor Model (FFM, often referred to as the Big5). Furthermore, related work [3, 43, 75] indicates that there are two more factors that have an influence on participants' privacy risk: *relationship type* (both relationship strength and equality of positions) and *preference scenario* (whether the active user's preference is in the minority or majority compared to others' preferences within the group).

Our results indicate that the following variables have a significant impact on the participants' privacy risk: two facets of personality (Extraversion and Agreeableness), preference scenario, and relationship type. These findings will inform the design of group explanation approaches in order to minimize privacy issues.

Therefore, we make the following key contributions:

- We investigate the relationship between factors identified in the literature that affect individuals' privacy risk: people's personality, their relationship type in the group, and preference scenario when explaining group recommendations.
- We conduct this user study with real groups of people (size=3). We implemented a web-based system where people could form a group and do the experiment.

The contribution of this study is published as a full paper in Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization [87].

4.2.1 Experimental Design

In this experiment, we investigate the following research question:

- How do people's personality, their relationship type in the group, and preference scenario affect their perceived privacy risk regarding group recommendation explanation?

Pre-Studies

Before starting the main study we needed to (1) verify that the exposure of personal information in the explanations actually does raise privacy risk perception in participants (pre-study 1) and (2) validate the instrument for measuring the privacy risk, which we adapted from related work (pre-study 2).

Pre-study 1: group explanation. Private information can fall under one or more of the following nine categories: location, medical, drug/alcohol, emotion, personal attacks, stereotyping, family or other associations, personal details, and personally identifiable information [9]. We selected the following subset which is relevant to the domain of tourism: location, drug/alcohol, emotion, personal details, and personally identifiable information. Some of this information is used in current tourism recommender systems, for example Mohamed et al. [82] use users' current location and emotion/mood, or Cheng et al. [11] consider user personally identifiable information (e.g., gender, age, race) to recommend personalized travel places to visit. In order to verify that the exposure of personal information in the explanations actually does raise privacy risk in participants, we ran a study. We asked ten colleagues from a computer science faculty to indicate how privacy sensitive each type of information, specifically in the context of an explanation given to the whole group, would be on a 5-point Likert scale ranging from 1 (non sensitive at all) to 5 (very sensitive). In addition, we provided an example from the explanation for each type of information. For instance, for the Drug/alcohol category (e.g., *you will love the Bulldog coffee-shop, a cannabis store*), for the Emotion category (e.g., *you are sad*), for the Personal details category (e.g., *your sexual orientation, LGBTQ+*), for the Personally Identifiable Information category (e.g., *your birth-date*), and for the Location category (e.g., *your current location*). The mean score was above 3 (out of 5) for all the types. This result suggests that the information used in the explanations is likely to provoke privacy risk.

Pre-study 2: establishing construct validity. Before we could measure the user's privacy risk regarding the presented group explanation, we needed to establish the validity of the instrument's items. We used confirmatory factor analysis (CFA), which can establish both the convergent (the question items are actually measuring a single construct) and discriminant validity (the question items are actually measuring different constructs). A rule of thumb for CFA is to have at least five participants per questionnaire item [58]. For 8 items, the minimum number of required sample size for our study is therefore estimated to be 40. The participants for this purpose were recruited using Prolific.⁴ We used results from 40 participants, after removing 5 participants who failed an attention check. The question items for measuring privacy risk were adapted for the purpose of our main study from previous instruments developed for measuring consumer information privacy in online contexts [62, 73]. We adopted a 5-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree). The values for both the CFI⁵ and TLI⁶ were above 0.95. The value of the Standardized Root Mean Square Residual (SRMR) was below 0.5. These values indicate a very good fit [58]. The question items with low squared loading values were removed from the final instrument (3 out of 8 items were removed).⁷ The remaining items were:

⁴www.prolific.co [July 2020].

⁵The Comparative Fit Index evaluates the model fit by analysing the discrepancy between the conjectured and the null model.

⁶The Tucker-Lewis index is preferable for smaller data samples and it indicates how much the conjectured model improves the fit relative to the null model.

⁷This decision is made on the basis of squared loading, using the recommended threshold of 0.50 [58].

- **P1)** The system disclosed, in this group explanation, information about me that I consider private.
- **P2)** All things considered, this group explanation would cause serious privacy problems.
- **P3)** To me, it is the most important thing to keep my privacy intact from the group members of the group I am in.
- **P4)** The system shows, in this group explanation, more information about me than I am comfortable with.
- **P5)** This group explanation is revealing too much personal information about me to the other group members.

4

Main study

In this section, we describe an online between-subjects study that investigates which factors influence the group members' privacy risk regarding the information disclosed in the presented group recommendation explanation.

We had two experimental manipulations (relationship type and preference scenario), an observed variable (personality), and a dependent variable (privacy risk). The *relationship type* variable takes the value of 1 for participants who are in a "loosely coupled heterogeneous group" (e.g., staff group including a manager), and 0 when they are in a "tightly coupled homogeneous" group (e.g., friend group). We controlled the *preference scenario* variable by setting its value to 1 for the member with minority preference (i.e., Carol in the scenario description in Section 4.2.1) and to 0 for the members that are in majority in terms of preferences (Bob and John from the scenario description).

We used the Big Five Inventory (BFI) to assess the *personality* on the five factors of Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism [49]. It is composed of 44 items with a 5-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree).

For the qualitative analysis, we distinguish between high and low for each trait. First, we aggregate all question items for each trait. Then after normalizing, we split the scores into two bins around the normative mean value for that trait as obtained in a large data-set of adult American Internet users [120]. With several question items, validated in the pre-study 2 (see Section 4.2.1), we measured the group members' *privacy risk* about a possible loss of privacy as a result of the group recommender system presenting an explanation to the group.

Procedure

In the previous section, we validated both the example explanation that we use in our user study and the question items for the main variable that we are measuring, privacy risk, in this user study. Here we introduce the online study we designed to evaluate people's privacy risk regarding a presented group explanation in the real groups.

We designed an online between-subjects experiment in which participants were randomly assigned to form either a) a tightly coupled homogeneous group or b) a loosely coupled heterogeneous group. The group size was set to three, similar to previous studies

of group recommendation [44, 74]. This was executed in two phases: 1) setting up groups, and 2) evaluating privacy risk.

Setting up groups. In this phase participants who see the advertisement and would like to participate in our study click on the provided link and are redirected to our “sign up” page. For each group, the experiment is initiated by one person, which we will refer to as an *inviter* as they are requested to invite the other two members.

- Step 1: While inviters sign up, they need to invite two other group members based on the group type we assign to them. For example for the “loosely coupled heterogeneous” type, they are requested to form a group where the position of members is unequal, the members are relatively estranged, and intercommunication is less frequent and less important (e.g., a staff group including a manager). By entering the potential group members’ names and emails, an invitation email is sent to the invitees.
- Step 2: When the two *invitees* get the invitation email they have a week to accept the invitation. Once both invitees accept the invitation, all group members get an email containing a link. This link which contains group members’ information redirects them to the second phase (evaluating privacy risk), corresponding to that group.

Evaluating privacy risk. When all the group members are redirected to the second phase, they go through the following steps:

- Step 0: Participants are shown a description of the scenario and the explanation for the recommended point of interest (POI) as can be seen in the following example.

The defined scenario:

Imagine that you and your group members have a plan to visit a place in Amsterdam together. A tourism app makes recommendations for your group based on all group members’ individual preferences. Carol has different tastes (preferences), compared to the other two group members. The recommended place is Carol’s favorite. The app has recommended the Bulldog coffee-shop to visit for your group which will explain why it made this recommendation for all of you as follows.

Names and ages will be adapted in the following explanations based on users’ inputs. Apart from that, it is the same for all the participants.

The presented explanation:

“The ”Bulldog coffeeshop” (cannabis store) has been recommended to your group since Carol will love it! The coffee-shop isn’t the primary preference of Bob and John, but they are okay with it. Their preferences will be taken into account in the next recommendations. Besides, Carol is feeling quite sad today, and we know that she really wants to visit the coffee-shop and won’t be talked out of it easily.

It’s a good recommendation geographically – it is close to all three of you. Carol is at Vondelpark, only a minute’s walk from the coffee-shop. Bob and John are at SoHo (LGBTQ+) bar, five minutes from the coffee-shop. You can all meet there in 10 minutes.

Since you’re all above 18 years in age, you can buy cannabis at the coffee-shop (Carol is 29, Bob is 28, and John is 35).”

4

- Step 1: Participants are asked to fill in some demographic-related questions as well as a set of questions to assess their personality traits.

- Step 2: Participants are asked to answer a set of survey questions related to their privacy risk of a shared explanation within the group in the defined scenario (the same questions validated in Section 4.2.1). We also include three attention check questions. At the end of the survey, participants are given the opportunity to freely express their opinions regarding the key factors that can influence their privacy preferences for a shared explanation in an open-ended question and the information they considered private.

Hypotheses

We investigate the relationship between factors identified in the literature that could affect privacy risk in the group, namely: *personality*, *relationship type*, and *preference scenario* and individual privacy risk. We developed a conceptual model to understand the relationship between those factors.

Figure 4.2 depicts the conceptual model, which includes a well-established connection (personality -> privacy risk) and two new connections that may also apply in the context of groups (relationship -> privacy risk and preference scenario -> privacy risk). We examine this conceptualization through the proposed theoretical lens of *privacy in groups*.




Figure 4.2: The conceptual model of disutility enhancers and reducers for perceived risk of disclosing personal information in a group explanation.

That leads to the following hypotheses:

- **H1a:** Agreeableness will influence individuals' privacy risk regarding the presented explanation to the group.
- **H1b:** Neuroticism will influence individuals' privacy risk regarding the presented explanation to the group.
- **H2:** The relationship type people have in the group will influence their privacy risk regarding the presented explanation to the group.
- **H3:** Preference scenario (having minority or majority preferences in the group) will influence individuals' privacy risk regarding the presented explanation to the group.

4.2.2 Results

To investigate the effects of different factors on the dependant variable, privacy risk, we built a structural equation model (SEM) upon the data collected with our questionnaire by using the R library Lavaan.⁸ All questionnaire items are modeled as ordinal variables. SEM is able to analyze the effects in an integrative structure where we can associate all the desirable effects.

The resulting SEM model (Figure 4.3) shows how type of *relationship*, *preference scenario* and *personality* influences privacy risk. Based on the final results we removed two question items (P1 and P2; which were validated in Section 4.2.1) with low squared loading values and three question items (P3, P4, and P5; see Section 4.2.1) remained valid to measure group privacy risk. The model has a good model fit: $\chi^2(205) = 340.423$, $p = .000$; root mean squared error of approximation (RMSEA) = .047; 90% CI : [0.026, 0.059], Comparative Fit Index (CFI) = .880, Tucker-Lewis Index (TLI) = .808.

⁸<http://lavaan.ugent.be/>, October 2020

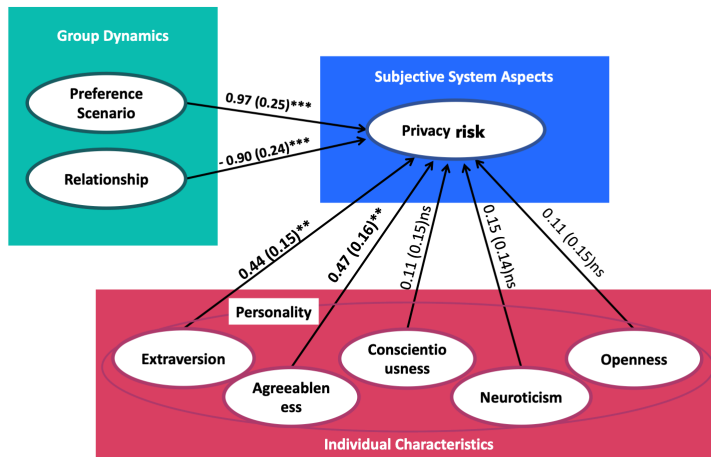


Figure 4.3: The structured equation modeling (SEM) results. Numbers on the arrows represents estimated coefficients (and standard error) of the effect. Significance levels: *** $p < .001$, ** $p < .01$, 'ns' $p > .05$.

Participants

To determine the required sample size, we performed a power analysis [28] of a medium-sized effect (0.5 SD) with a power of 85% in a between-subjects experiment. It showed that a minimum of 100 participants are needed in total. This was inline with the suggested minimum sample size for SEM in Knijnenburg and Willemsen [58].

The participants for this study were a convenience sample recruited through university networks. 114 participants (38 groups of 3 people) voluntarily joined our study (Age: Mean = 31.8, SD = 7.7; Gender: Female = 47%, Male = 53%). Half of the participants were assigned to form a loosely coupled heterogeneous group (19 groups) and the other half a tightly coupled homogeneous group (19 groups). By design,⁹ among those one-third of participants (38 participants) were assigned to have minority preferences in the group and two third majority preferences (76 participants). All responses were included in the data analysis due to successful attention checks.

H1. Effects of Personality

Here we discuss the effects of the two personality traits we hypothesised would have an effect on privacy risk:

H1a. Agreeableness. We found that the Agreeableness trait in our participants has significant effect on their privacy risk ($p < .01$). The positive sign (coefficient=0.47) indicates that people who scored high on Agreeableness perceived higher privacy risk rather than people who scored low on this trait. Thus, we can accept hypothesis H1a: Agreeableness will influence individuals' privacy risk regarding the presented explanation to the group.

⁹Recall that in each group, one participant preferred the POI and this was in contrast with the preferences of the other two group members who constituted the majority.

H1b. Neuroticism. We found no significant effect of participants' score on the Neuroticism trait on the participants' privacy risk. We argue that one possible reason that we did not find an effect could be the distribution of scores on this trait in our sample. Only 19 participants scored high on this trait in comparison to 95 participants with a low score on this trait.

H2. Effects of Relationship

We found that people's relationship type within the group has a significant effect on their privacy risk ($p < .001$). Specifically, the negative sign (coefficient= -0.90) indicates that participants in a tightly coupled homogeneous group (e.g., friend group) perceived a lower privacy risk, compared to participants in a loosely coupled heterogeneous group (e.g., staff group). Thus, we can accept hypothesis H2: The relationship type people have in the group will influence their privacy risk regarding the presented explanation to the group.

H3. Effects of Preference Scenario

We found that having the *minority* or *majority* preferences in the group has a significant effect on people's privacy risk ($p < .001$). Specifically, the results are consistent with conformity: the positive sign (coefficient= 0.97) indicates that people having minority preferences perceived higher privacy risk rather than people having majority preferences. Thus, we can accept hypothesis H3: Preference scenario (having minority or majority preferences in the group) will influence individuals' privacy risk regarding the presented explanation to the group.

Post-hoc Analysis

The results from related work for the trait of Extraversion were weak and inconsistent. Given that we included all the personality traits in our model, we are able to report the results for Extraversion here for further comparison. Extraversion reflects people orientation and pleasure in social interactions. Descriptions of this trait include being talkative, bold, assertive, sociable, and demonstrative [36]. We found that the Extraversion trait in our participants has a significant effect on their privacy risk ($p < .01$). The positive sign (coefficient= 0.44) indicates that people who scored high on Extraversion perceived higher privacy risk rather than people who scored low on this trait. We argue that a possible reason that we found a significant and strong effect could be the distribution of scores on this trait in our sample. We had a similar number of participants who scored high and low on this trait.

Qualitative Feedback

We asked the participants to motivate their responses. In this section, we analyse their comments to better understand whether the three factors (personality, relationship type, and preference scenario) influenced privacy risk differently for the five information types (Location, Drugs, Emotion, Personally identifiable Information, and Personal Details). The users' feedback was analysed using closed/fix coding [71]. They were coded based on one or more information categories (among the five information categories) users mentioned in their comments. For example, we coded the comment: "*The description indicates both that I would love a **cannabis cafe**, and that I am **very sad**. I consider both statements to be way too personal and private!*" for the category of *Drugs* and the category of *Emotion*.

Then we divided these categories based on our three main factors: their relationship type, preference scenario, and their personality. All participants' comments were considered in this analysis. Following we describe the results in detail.

Location. *Personality.* 31% (17 out of 54) who scored high on Extraversion, and 31% (6 out of 19) who scored high on Neuroticism expressed privacy risk in their comments regarding disclosing their current location. Also, 25% (18 out of 72), who scored high on Agreeableness showed privacy risk about revealing their current location. Conversely, only a few participants, about 5% of participants who scored low on these three traits, showed privacy risk regarding this type of information.

Relationship type. 21% (12 out of 57) of participants who were in a loosely coupled group expressed privacy risk about revealing their current location. From the participants in a tightly coupled homogeneous group, only half of this number of participants (10%, 6 out of 57) expressed privacy risk for disclosing their current location to the group.

Preference scenario. 21% (8 out of 38) of participants who had minority preference within the group expressed privacy risk in their comments regarding disclosing their current location. Fewer participants (16%, 12 out of 76), who had majority preferences, expressed privacy risk for this type of information.

Drug/alcohol. Overall, participants expressed less privacy risk of disclosing this type of information.

Personality. The highest number belongs to participants who scored high on Extraversion: 15% (8 out of 54). A comparable number of participants who scored high on Agreeableness and Neuroticism expressed privacy risk of disclosing this type of information, both 10%.

Relationship type. 17% (10 out of 57) of participants who were in a loosely coupled group expressed privacy risk about revealing drug/alcohol information. From the participants in a tightly coupled homogeneous group, only 2% (1 out of 57) expressed privacy risk for disclosing this type of information to the group.

Preference scenario. 18% (7 out of 38) of participants who had minority preference within the group showed privacy risk in their comments about disclosing drug/alcohol information. A small proportion of participants about 8% (6 out of 76) of participants, who had majority preferences, showed privacy risk for this type of information.

Emotion. *Personality.* Participants perceived more risk of disclosing their emotional state in the group. The highest number belongs to participants who scored high on Neuroticism: 47% (9 participants out of 19). Besides, 37% (20 out of 54) of participants who scored high on Extraversion perceived risk for this type of information in their comments. Fewer participants about 18% (13 participants out of 72), who scored high on Agreeableness, expressed this privacy risk.

Relationship type. 28% (16 out of 57) of participants who were in a loosely coupled group expressed privacy risk about revealing their emotion. From the participants in a tightly coupled homogeneous group, about 17% (10 out of 57) of the participants expressed privacy risk for disclosing this type of information to the group.

Preference scenario. 34% (13 out of 38) of participants who had minority preference within the group expressed privacy risk in their comments regarding disclosing their emotional state. Fewer participants about 10% (8 out of 76), who had majority preference within the group expressed privacy risk for this type of information.

Personally identifiable information (age). *Personality.* 30% (16 out of 54) of participants who scored high on Extraversion perceived risk about disclosing their age within the group. Besides, 25% (18 out of 72) of participants who scored high on Agreeableness perceived risk for this type of information in their comments. From participants who scored high on Neuroticism, 19% (4 out of 19) stated this privacy risk.

Relationship type. 26% (15 out of 57) of participants who were in a loosely coupled group perceived risk about revealing their age. From the participants in a tightly coupled homogeneous group, about 10% (6 out of 57) perceived privacy risk for disclosing this type of information to the group.

Preference scenario. 24% (9 out of 38) of participants who had minority preference within the group expressed privacy risk in their comments regarding disclosing this type of information. Fewer participants, about 14% (11 out of 76), who had majority preferences, expressed privacy risk for this type of information.

Personal details (LGBTQ+). *Personality.* The highest number of participants who perceived privacy risk about disclosing this type of information to the group were the ones who scored high on Agreeableness: 26% (19 out of 72). For participants who scored high on the two other traits, Extraversion and Neuroticism, fewer participants perceived privacy risk for this type of information in their comments (10%).

Relationship type. 28% (16 out of 57) of participants who were in a loosely coupled group showed privacy risk about revealing their sexual orientation. From the participants in a tightly coupled homogeneous group, fewer participants, about 5% (3 out of 57), showed privacy risk for disclosing this type of information to the group.

Preference scenario. 31% (12 out of 38) of participants who had minority preference within the group showed privacy risk in their comments regarding disclosing this type of information. Fewer participants, about 10% (8 out of 76), who had majority preferences, showed privacy risk for this type of information.

4.2.3 Discussion

Following this, we look closer at our results and their implications.

Personality. We found that participants who scored high on Agreeableness or Extraversion perceived more privacy risk of information disclosure. This was further supported by the qualitative comments from participants. The comments indicate that participants who scored high on Agreeableness were concerned more about their location, age, and sexual orientation information than about other types of information (about 19% of these participants). In contrast, the highest privacy risk for participants who scored high on Extraversion was about their emotional information (about 37% of these participants). This was similar for participants who scored high on Neuroticism (about 47% of these participants), who showed more privacy risk about their emotional information. As a guideline

for designing explanations, we should adapt which information is disclosed depending on the personalities in the group. Different personality traits varied in terms of which information they found sensitive.

Relationship type. The highest number of participants, who perceived privacy risk about sexual orientation and emotional information were in a loosely coupled heterogeneous group (about 28% of these participants). The highest number of participants, who were in a tightly coupled homogeneous group perceived privacy risk about only emotional information (about 17% of these participants). We should adapt to the loosely coupled heterogeneous group for all five information types.

4

Preference scenario. The highest number of participants who perceived privacy risk about emotional information had minority preferences (about 35% of these participants). This was about location information for the participants who had majority preferences (about 16% of these participants). Minority preferences matter a little for all information types, but in particular for emotion. We should adapt to people with minority preferences in particular when disclosing emotion.

Information type. Among the five types of information we included in the explanation, the highest number of participants perceived privacy risk for the emotional information. Information regarding participants' age came second with regards to the privacy risk. Surprisingly, Drugs appear to be the least important to adapt to (18% max). This might be related to the cultural background of our participants. As our participants mainly live in the Netherlands, maybe in the Netherlands people are less sensitive in disclosing this type of information. In the future, it would be interesting to study the relationship between the nationality/where participants live and their privacy risk for different types of information.

Setting up groups. Recruitment of groups participants is a challenge when aiming to control for the group type. The challenge increases when recruiting heterogeneous, loosely coupled groups, in particular with a leader. A recommendation for future studies is to first ask the participants from a "higher" position, rather than to recruit organically or to request participants in "lower" positions to recruit others. We received feedback from several participants that it is difficult for them to ask a person in a higher position (e.g., their boss) to form a group with them.

Limitations

In this section, we discuss the limitations of our study.

Firstly, we measured participants' privacy risk for a hypothetical scenario rather than their actual preferences. This might cause people not to be able to imagine the situation very well. Although we used actual groups and participants' answer to the open-ended question shows their high engagement in the study, asking participants to imagine sharing their information still might lead to different results than actually sharing it.

Secondly, we defined the scenario in such a way that we expected to maximize the privacy risk. In future studies, we plan to study the effect on different information types

in more detail. Besides, in our study we studied preference in relation to a single POI which was sensitive due to being a coffee shop. Different results might be found for other types of preference scenario.

In addition, participants were recruited from universities worldwide through our professional networks. This sample may not be representative of the general population. For example, an effect for the personality trait of Neuroticism might be found in a sample that controls for the balance of high and low scores on this trait. Or the sensitive information is probably less sensitive culturally for the majority of the members of this sample.

Our work could also benefit from the larger sample size. In this study, we considered a bare minimum for relatively simple SEM models (100 observations). Besides, we relied on the independence of the data points in our model.

Finally, this study was conducted in the context of recommendations for tourism. This domain was suitable for studying group recommendations, as it is relatable for many participants. However, the results may differ in domains where preferences are less subjective in nature or differ in terms of their level of investment or risk [126].

In this section, we studied some factors identified in the literature that influence individual's perceived privacy risk of disclosing personal information, i.e., group members' personality, the type of relationship they have in the group. Also on the basis of the previous experiment (the privacy preferences experiment), i.e., preference scenario (whether their preferences are aligned or not aligned with the preferences of the majority in the group). We studied these factors when disclosing several types of personal information (i.e., location, drug/alcohol, emotion, personal details, and personally identifiable information) in a single group explanation. The results helped us to see that these factors generally affect individual privacy risk. As we found significant effects of the identified factors on users' privacy risk perception, in the next section, we study the effect of these factors on people's actual information disclosure (see Section 4.3). We expect to see the opposite effect of factors on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure).

4.3 Experiment 3: Disclosing Location and Emotion Information

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 [124]. Many studies have demonstrated the benefits of adding explanations to automated recommendations (e.g., [42, 117]). Previous research in this area has focused on explaining individual recommendations [42, 117]. When explaining recommendations to a group of users, it is challenging to recommend an item to a group that satisfies all group members simultaneously [2]. 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 [75], e.g., consider the explanation “*Alice is feeling sad today, and she really wants to visit this place*”.

4

In the privacy factors experiment (see Section 4.2), we found three factors that influence the perceived privacy risk of information disclosure in the tourism group recommendation context, namely, group members’ *personality* (modeled using the Five-Factor Model [14]), *preference scenario* (whether the active user’s preferences align with the majority in the group or not), and the type of *relationship* (the relationship strength between group members and equality of their positions). However, in that work, we looked at disclosing five kinds of different personal information (i.e., location, drug/alcohol, emotion, personal details, and personally identifiable information) in a single group explanation. So this needs further investigation to find out which of these personal information types should be tailored for different personalities or group composition. Qualitative analysis from user comments suggested that for example people with different personality traits are concerned differently regarding different types of personal information (the privacy factors experiment, see Section 4.2). 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. [80]. In this experiment (see Section 4.3), we consider the tourism domain and consider the specific personal information types individually rather than all in a single explanation. Namely we study *location* and *emotion*, which are most used in current tourism recommender systems (e.g., [82]). Another main distinction to our previous work (the privacy factors experiment) is that we looked at participants’ actual information disclosure rather than their perceived privacy risk of information disclosure.

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 perceive a larger privacy risk regarding the disclosure of emotion-related compared to location-related information. In contrast to previous research, we find no evidence that personality traits or preference scenarios affect privacy risk. 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., 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. Therefore, we make the following key contributions:

- We study how individual differences (i.e., personality) and group compositions (i.e., preference scenarios) affect people’s privacy risk of location-related and emotion-related information disclosure when explaining group recommendations to their group.
- To create realistic scenarios of group decision-making where users can control the amount of information disclosed in the group and have iterative interaction between group members, we developed TouryBot. This web-based chat-bot agent generates natural language explanations to help group members explain their suggestions to the group in the tourism domain.

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 contribution of this study is published as a full paper in Proceedings of the 32nd ACM Conference on Hypertext and Social Media [88].

4.3.1 Experimental Design

In this section, we describe an online between-subjects study that investigates how personality and preference scenario relate to individuals’ privacy risk 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.

Specifically, we investigate the following research questions:

1. How do personality and preference scenario affect people’s location-related privacy risk in explaining group recommendations to their group?
2. How do personality and preference scenario affect people’s emotion-related privacy risk in explaining group recommendations to their group?

Study Platform

To answer the research question we implemented a web-based chat-bot that we call TouryBot. For the UI we used a client in java (Vaadin AI Chat)¹⁰ and implemented in Vaadin framework.¹¹ The backend is written in python. SQLite was used for logging user inter-

¹⁰<https://github.com/alejandrod/uaadin-ai-chat>, retrieved March 2021.

¹¹An open platform for building web apps in Java (<https://vaadin.com/>), retrieved March 2021.

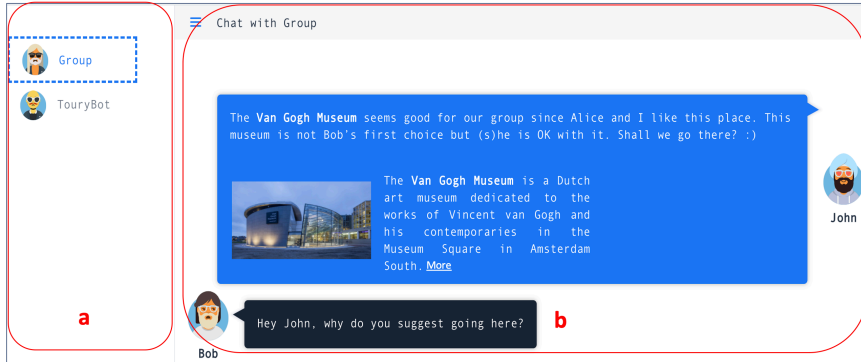


Figure 4.4: The chat in the majority preference scenario between an active user and his group. a) indicates two ongoing chats, one with a chat-bot and the group, b) indicates the active user (John) shares his preference (the Van Gogh museum in this example) with his two group members (the two other group members, Bob and Alice, are hypothetical) in a group chat.

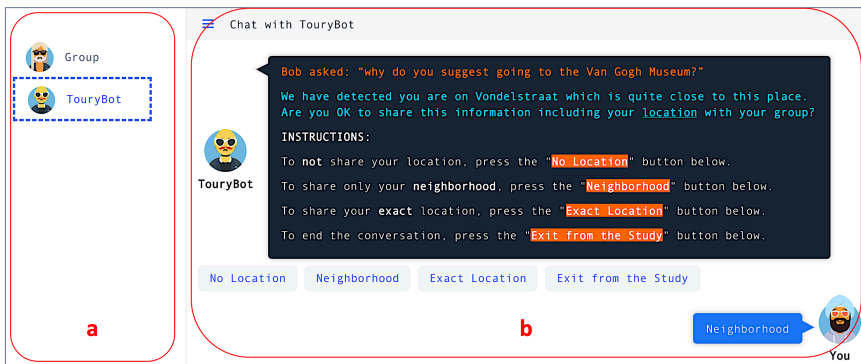


Figure 4.5: The chat in the majority preference scenario between the chat-bot and an active user. a) indicates two ongoing chats, one with a chat-bot and the group, b) indicates an ongoing chat with a chat-bot. Here the user can indicate the level of location information they want to share to convince the other group member (Bob) to visit the suggested POI.

actions in the task. TouryBot includes two chat windows, one for the chat between the system bot and individual members (see Figure 4.5), and the other for the chat with the Group (see Figure 4.4). Users can seamlessly switch between the two chats to add system-generated recommendations and explanations to their discussions with other group members.

Measures

Inspired by our previous experiment (the privacy factors experiment, see Section 4.2), this 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 risk as dependent variables.

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* [49]. 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.

Dependent variables

Location-related privacy risk (ordinal). We used three different levels of granularity for location-related information to measure the group members' *privacy risk* 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 risk (ordinal). Similarly, we used three different levels of granularity for emotion-related information to measure the group members' *privacy risk* 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.

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.

4

Materials

Emotion content

This study concerns users' willingness to disclose *emotion*- and *location*-related information, which are among the five personal information types used in our previous experiment (the privacy factors experiment, see Section 4.2). These two information types have been used in current tourism recommended systems; e.g., Mohamed et al. [82] 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 risk in participants to be included in the explanation. Note that no pre-study was conducted for *location-related information* as privacy risk about disclosing current location has been studied extensively (e.g., [132]).

Ekman and Friesen [24] 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*) [106].

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.

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*.¹² By using participants’ real preferences, we aimed to increase the likelihood of a more realistic situation for users to imagine.

Procedure




Figure 4.6: 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.

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 4.6a). 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 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 4.3.1. 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 (Figures 4.6b). 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

¹²<https://developer.foursquare.com/>, retrieved February 2021.

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 4.3.1. 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 4.4, Bob (the hypothetical group member) asks why John (the active user) suggests that POI. And as can be seen in Figure 4.5, 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 4.5b).

Personality (Figure 4.6c). After completing the scenario, participants filled in the BFI for assessment of their personality traits.

Participants

To determine the required sample size, we performed a power analysis [19] 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*.¹³ This platform has shown to be an effective and reliable choice for

¹³<https://www.prolific.co>

running relatively complex and time-consuming studies, e.g., for interactive information retrieval [138]. 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).

Hypotheses

In Section 2.4, we discussed relevant literature on antecedents of privacy risk and disclosure benefit and ultimately their effect on final personal information disclosure in group recommendation context. Based on that we formulate the hypotheses that lead this work as follows:

- **H1.a)** Extraversion affects location-related information disclosure.
- **H1.b)** Neuroticism affects location-related information disclosure.
- **H1.c)** Agreeableness affects location-related information disclosure.
- **H1.d)** Participants whose preferences are in the minority disclose less location-related information compared to participants whose preferences are in the majority.
- **H2.a)** Extraversion affects emotion-related information disclosure.
- **H2.b)** Neuroticism affects emotion-related information disclosure.
- **H2.c)** Agreeableness affects emotion-related information disclosure.
- **H2.d)** Participants whose preferences are in the minority disclose less emotion-related information compared to participants whose preferences are in the majority.

Statistical Analyses

To test our hypotheses, we performed two *ordinal logistic regression* (OLR) analyses [40] (i.e., one to predict location-related and the other for emotion-related information disclosure) with *preference scenarios* and the five personality scales *extraversion*, *agreeableness*, *neuroticism*, *openness*, and *conscientiousness* as independent variables.¹⁴ We corrected for multiple hypothesis testing by lowering the significance threshold to $\frac{0.05}{8} = 0.00625$ (i.e., applying a *Bonferroni correction* [89]).

4.3.2 Results

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

¹⁴Although our hypotheses concerned only the first three of the five personality scales, we added *openness* and *conscientiousness* as covariates to the models to account for potential confounding factors.

Hypothesis Tests

Tables 4.4 and 4.5 show the results from the OLR analyses regarding location information and emotion information respectively. We found *no evidence* in favor of any of our eight hypotheses (**H1a** - **H2d**; all $p > 0.00625$; see also Section 4.3.1). 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 risk (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 risk (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**).

Table 4.4: Results of location-related information disclosure from an ordinal logistic regression (OLR) analyses in a group explanation as a dependent variable (DV). 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 4.3.1). However, no factors were statistically significant after correcting for multiple testing (see Section 4.3.2).

DV: Location-Related Privacy Risk				
Hyp.	Factor	β	p	OR [95% CI]
-	Intercept 1 2	-1.11	0.38	
-	Intercept 2 3	1.66	0.19	
H1a	extr	-0.07	0.65	0.94[0.70, 1.25]
H1b	agr	-0.21	0.26	0.81[0.56, 1.17]
H1c	neur	0.05	0.71	1.06[0.80, 1.40]
-	open	-0.17	0.31	0.84[0.60, 1.18]
-	cons	0.04	0.78	1.04[0.77, 1.42]
H1d	pref	-0.09	0.66	0.91[0.60, 1.38]

In sum, based on the OLR results, we cannot reject any of the null hypotheses opposing the hypotheses we aimed to test (see Section 4.3.1). 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 risk) may be absent or much smaller than previously anticipated in this context.

Table 4.5: Results of emotion-related information disclosure from an ordinal logistic regression (OLR) analyses in a group explanation as a dependent variable (DV). 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 4.3.1). However, no factors were statistically significant after correcting for multiple testing (see Section 4.3.2).

DV: Emotion-Related Privacy Risk				
Hyp.	Factor	β	p	OR [95% CI]
-	Intercept 1 2	-1.83	0.12	
-	Intercept 2 3	-0.10	0.93	
H2a	extr	-0.19	0.18	0.83[0.63, 1.09]
H2b	agr	-0.08	0.66	0.93[0.65, 1.31]
H2c	neur	-0.14	0.31	0.87[0.67, 1.14]
-	open	-0.11	0.51	0.90[0.65, 1.24]
-	cons	0.27	0.06	1.32[0.99, 1.75]
H2d	pref	-0.42	0.03	0.66[0.45, 0.97]

Exploratory Findings

We would expect to find an effect of personality and preference scenario on participants' information disclosure to their group, however, surprisingly we did not find any effect. In this section, we present several exploratory findings that may help to explain this surprising results from the hypothesis tests.

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 risk ($p = [0.47, 0.90]$; results of ordinal logistic regressions).

Partial disclosure. Table 4.6 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.

Information type. In line with previous research [80], we found that privacy risk varies depending on information types, with significantly larger risk of disclosure 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 4.6 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.

Table 4.6: Number (and percentage) of participants across privacy risk of location (left) and emotion (right) information disclosure.

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

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.

Qualitative Feedback

We asked the participants to motivate their responses. In this section, we analyse their comments to better understand why they disclose more or less personal information to the group.

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.”*

4.3.3 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 risk in a realistic chat-bot scenario (see Section 4.3.1). Our results contain no evidence for any of these effects (see Section 4.3.2). 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 the previous experiment that suggests that the preference scenario, as well as personality traits, do affect users’ privacy risk (the privacy factors experiment, see Section 4.2). 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 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”. For future work, it would be helpful to look at whether the “convincing mindset” shelters these effects.

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 our previous experiment (the privacy factors experiment, see Section 4.2) that reported effects of preference scenario and personality traits on privacy risk, where particularly a sensitive POI was used (e.g., a cannabis store). The arguably lower overall privacy sensitivity in our study might have diminished these effects, e.g., causing people having minority preferences to feel less risk 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 perceived risk regarding disclosing these types of information in a group explanation (the privacy factors experiment, see Section 4.2). Furthermore, in our study especially 40% chose partial disclosure of each information type (e.g., neighborhood location, approximate emotion) which shows people do perceive privacy risk 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. Our study design diminished previously demonstrated effects (i.e., the factors identified in the privacy factors experiment, did not affect people's disclosure behavior.), which might be important for designing chat-bots in such a group recommendation scenario. This can help to avoid nudging people into some "convincing" mindset as in this context they might disclose more personal information than they are comfortable with.

Limitations and Future Work

General privacy concern. In this study, we did not find an effect of personality traits and preference scenario on privacy risk. There might be an additional mediating factor that affects participants' privacy risk. For example, it would be beneficial in future work to also measure general privacy concern [61] to see if it mediates privacy risk of disclosing 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 perceive higher privacy risk. 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 risk in more details.

User control. To be able to study participants' privacy risk 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.

We conducted three experiments discussed in this chapter to answer which factors we should model in the group to consider group members' privacy risk of information disclosure. The following section (see Section 4.4) summarizes our findings of the experiments in this chapter.

4.4 Chapter Conclusions

In a group recommendation/decision context, there is information that we can present to people in a group to help them reach a consensus on the recommended items. Adding more detailed information in situations when there is disagreement in the group seems more beneficial. However, this leads to an increase in individuals' privacy risk perception. This chapter covered three experiments looking at the factors influencing individuals' privacy risk perception of information disclosure and, ultimately, their disclosure behavior in the group. The privacy preferences experiment investigated which information people would like to disclose or not disclose in explanations for group recommendations. The results suggested individual differences regarding privacy risk perception of information disclosure. To understand what contributes to these individual differences, the privacy factors experiment investigated the relationship between factors identified in the literature and on the basis of the previous experiment that affect individuals' privacy risk. As we found significant effects of these factors on users' privacy risk perception, we studied the effect of these factors on users' actual disclosure in the information disclosure experiment. We expected to see the opposite effect of factors identified in the privacy factors experiment, on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure). However, neither the personality traits nor the preference scenario affected people's actual disclosure.

4

Experiment 1: privacy preferences. We presented a framework adapted to users' privacy preferences to generate natural language explanations for groups, aiming to evaluate which information to disclose in explanations for group recommendations. We compared these privacy preferences for different low consensus scenarios, where either the active user or their acquaintances did not get their preferred item (when the group has a disagreement on the recommended item), and a high consensus scenario where everyone got their desired item (when the group has an agreement on the recommended item).

We found that people use more privacy options in both low consensus scenarios than in high consensus scenarios. They also avoid selecting the combination of name and a strong opinion of a group member(s). Moreover, users' responses to the evaluation questions encouraged findings that our provided privacy control options were correct and had enough choices.

In line with previous work, [62], our findings suggest that there may be some individual differences in the levels of participants' concern about being singled out for having different preferences. Whether we can predict which factors (e.g., personality, relationship, etc.) contribute to these differences led our following study to investigate this further.

Experiment 2: privacy factors. In this study, we investigated and found an effect of three factors influencing individuals' privacy risks regarding an explanation of a tourism group recommendation: relationship, preference scenario, and personality. *Relationship type* has a substantial effect on privacy concerns.

The results showed that participants in a tightly coupled homogeneous group (the relationship strength between group members is high, and their positions in the group are equal) perceived lower privacy risk than in a loosely coupled heterogeneous group (the relationship strength between group members is low and their positions in the group are not

equal). Besides, the *preference scenario* (whether group members' preferences are aligned or not aligned with the preferences of the majority in the group) also has a substantial effect on privacy risk. Participants in the minority preference-wise perceived higher privacy risk than group members in the majority preference-wise. For *personality*, we found the expected effect of the personality trait of Agreeableness on privacy risk. However, we did not find the hypothesized impact of the trait of Neuroticism. Additionally, we found an effect on the trait of Extraversion. Both Agreeableness and Extraversion positively impacted privacy risk, meaning that an individual that shows high levels of Agreeableness or Extraversion would perceive higher privacy risk.

Moreover, the participants' comments suggested individual differences regarding which information to disclose in relation to the three factors. This leads the following study to investigate these factors regarding each individual information type.

Experiment 3: information disclosure. We presented a user study investigating the effect of three personality traits and preference scenarios on users' actual information disclosure. We found that information disclosure varies depending on information types, with significantly less disclosure for *emotion-related* than *location-related* information. Although we have expected to see the opposite effect of factors, identified in the privacy factors experiment, on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure), neither the personality traits nor the preference scenario affected people's disclosure behavior. One explanation for why we did not obtain the expected results could have been that our *task design* nudged participants into a "convincing mindset". Our tasks required participants to convince other group members, and overall, people disclosed more in a convincing mind-set. However, this could have caused the "convincing mindset" overshadows the effects of previously identified factors on actual disclosure, e.g., group configuration and relationship. Besides, several studies show that when people decide on personal information disclosure, they trade off the anticipated *benefits* (i.e., the extent to which users believe disclosing their personal information to their group members is beneficial for the group decision or their negotiation position) with disclosure *risks* (i.e., the expectation of losses associated with the disclosure of personal information to the group). This remains for the next chapter to investigate this trade-off.

Wrap up

- When the recommended item is aligned with the minority's preferences in the group rather than the majority's preferences, disclosing the identity (i.e., name) of the people with the minority preferences in the group, together with their strong opinions, should be avoided. Instead, the explanation is better to be made anonymous. For example, *some of the group members want to visit this place and won't be talked out of it easily*.
- Our results suggested that information disclosure varies depending on information types, with significantly larger disclosure for *location-related* than *emotion-related* information.

- We found contradictory results between the factors influencing people's privacy risk of information disclosure and their actual disclosure behavior. One explanation for why we did not obtain the expected results could have been that our task design nudged participants into a "convincing mindset". Besides, to predict group members' actual disclosure behavior, we should consider other intermediate factors that ultimately affect individuals' disclosure behavior, not only the final behavior. For example, measuring group members' perceived privacy risk and disclosure benefit.

Based on these findings, in the next chapter, we focus more on the effect of task design on actual information disclosure in the group, as well as trade-offs between privacy risk and disclosure benefit and their effects on actual information disclosure.

5

Antecedents of Information Disclosure

5

To design textual explanations for group recommendations, in Chapter 3, we started by determining what information needed to be conveyed to people in a group explanation for the recommended items. User comments highlighted the need for protecting certain types of information when presenting an explanation to the group, i.e., group members' ratings of items. This leads us to investigate the factors that we should model in the group to consider group members' perceived privacy risk of information disclosure. In Chapter 4, we investigated the important factors that influence group members' perceived privacy risk (i.e., expectation of losses associated with the disclosure of personal information in the group) regarding the generated explanations. In the privacy factors experiment (see Section 4.2), we found that group members' personalities (using the 'Big Five' personality traits), their preference scenarios (i.e., whether their preferences are aligned or not aligned with the preferences of the majority in the group), and the type of relationship they have in the group (i.e., loosely coupled heterogeneous like a colleagues group, versus tightly coupled homogeneous like a friends group) have a strong influence on people's privacy perception. In a follow-up experiment (the information disclosure experiment, see Section 4.3), we investigated the effects of these factors on people's disclosure behavior (how they choose, if any, among the certain types of personal information to share with their group members). Although we have expected to see the opposite effect of factors on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure), neither the personality traits nor the preference scenario affected people's disclosure behavior. One explanation for why we did not obtain the expected results in the information disclosure experiment could have been that our task design nudged participants into a "convincing mindset". To understand this better, the Task design experiment (see Section 5.1) investigates the effect of task design on people's actual disclosure. The results revealed that task design (i.e., the pressure on users to convince the group) affected participants' disclosure decisions (for certain types of sensitive information). The influence of task on actual disclosure leads to the idea that disclosure benefit and privacy risk of information disclosure could cancel each other out in actual disclosure decision and result in smaller effects of the factors mentioned earlier (i.e., personality, etc.) on participants'

actual information disclosure. This leads us to the third research question, which investigates how people make the trade-off between disclosure benefit versus privacy risk to deciding on information disclosure in a group explanation (RQ3). Disclosure benefit in our context refer to the extent to which users believe disclosing their personal information to their group members is beneficial for the group decision or for their own negotiation position within the group. In fact, several studies show that when people decide on personal information disclosure, they trade off the anticipated benefits with the risks of disclosure. In the benefit vs. risk experiment (see Section 5.2), we further unpack participants' disclosure decision by measuring their perceived privacy risk and disclosure benefit. We find that these factors mediate the effect of the previously mentioned factors (i.e., personality, etc.) on participants' actual disclosure behavior.

This chapter is based on part of a conference paper and a submitted journal paper:

- Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. Exploring user concerns about disclosing location and emotion information in group recommendations. In Proceedings of the 32nd ACM Conference on Hypertext and Social Media, pages 155–164, 2021
- Shabnam Najafian, Geoff Musick, Bart Knijnenburg, and Nava Tintarev. How do People Make Decisions in Disclosing Personal Information in Tourism Group Recommendations in Competitive versus Cooperative Conditions? Modeling and User-Adapted Interaction (UMUAI) Journal, 2022, (under review)

5.1 Experiment 1: Influence of Task Design on Information Disclosure

This thesis aims to study what makes good explanations for group recommendations. Our findings highlighted that people might be sensitive to disclose some of their information in a group, e.g., their preferences (see Chapter 3). This led us to investigate the factors that we should model in the group to consider group members' privacy risk of information disclosure (see Chapter 4).

In the privacy factors experiment (see Section 4.2), we found three factors that influence the group members' perceived privacy risk in the tourism group recommendation context, namely, group members' *personality* (modeled using the Five-Factor Model [14]), *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). In the information disclosure experiment, we aimed at further investigation to find out how these factors affect group members' actual disclosure behavior (see Section 4.3). Although we have expected to see the opposite effect of factors on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure), neither the personality traits nor the preference scenario affected people's disclosure behavior.

The information disclosure experiment required participants to *convince/persuade* other group members. Therefore, one explanation for why we did not obtain the expected results could have been that our task design nudged participants into a "convincing mindset". So in this study, we investigate the effect of task design on people's actual disclosure.

To answer the above question, we designed a user study where participants receive recommended *point of interests* (POIs) from the group recommender in both majority and minority preference scenarios and are exposed to either persuasive or non-persuasive (descriptive) task design. To be able to investigate the effect of task design we used the exact same experimental setup as the previous study (the information disclosure experiment), using TouryBot, a chat-bot agent that supports the natural dynamics of group decision-making.

Our results revealed that task design had a strong effect on participants' emotion-related disclosure behavior. In particular, the *nudging* of users to convince the group can partly explain such lack of evidence. 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 – that should be taken into account when designing chatbots and similar tools for decision-making in group recommendations contexts.

Therefore, we make the following key contributions:

- We study the effect of task design (whether group members were instructed to convince other group members of their opinion or not) on participants' information disclosure in the group.
- To create realistic scenarios of group decision-making where users can control the amount of information disclosed in the group and have iterative interaction between group members, we developed TouryBot. This web-based chatbot agent generates natural language explanations to help group members explain their suggestions to the group in the tourism domain.

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 contribution of this study is published as a full paper in Proceedings of the 32nd ACM Conference on Hypertext and Social Media [88].

5.1.1 Experimental Design

As this experimental setup is the same as in the experimental setup described for information disclosure experiment (see Section 4.3), we do not repeat the details in this section but we only mention the changes. The information disclosure experiment exposed users to one of two tasks: either (1) majority preference scenario where the active user need to convince other group members to *accept* visiting the suggested POI or (2) minority preference scenario where the active user need to convince other group members to *skip* the suggested place. Both tasks thus required participants to *convince* other group members. To investigate if the task design of convincing other group members had an effect on our study regarding actual disclosure, 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 4.3.1, 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 information disclosure experiment (see Section 4.3.1). A required sample size of 180 additional participants was computed in a simulation study beforehand. 179 participants remained after removing those who failed the attention checks. We added this additionally obtained data to our data set of 362 participants from the previous study, 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 ordinal logistic regression (OLR) analyses [40] again with *convincing* as an additional factor. Otherwise, the OLR analyses were performed in the same way as described in Section 4.3.1.

Hypotheses

Based on what we discussed previously we formulate the hypotheses that lead this work as follows:

- **H1)** Task design affects location-related information disclosure.
- **H2)** Task design affects emotion-related information disclosure.

5.1.2 Results

In this section, we discuss the outcomes of the hypothesis tests we conducted.

Hypothesis Tests

Whereas *convincing* did not have an effect regarding location-related information disclosure ($\beta = -0.25$, $p = 0.17$, $OR = 0.78$ with 95% $CI[0.55, 1.11]$), it did affect emotion-related information disclosure ($\beta = -0.97$, $p < 0.001$, $OR = 0.38$ 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 have to convince other group members. Although we did not find the same effect regarding location-related information disclosure, the trend there went in the same direction as for emotion-related information disclosure. It could thus be that the convincing aspect affected location-related information disclosure 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 4.3.2 (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.

5.1.3 Discussion

Task design. Our task design had an effect on emotion-related information disclosure (see Section 5.1.2), which could be one of the reasons for getting contradictory results between two previous experiments, namely the privacy factors vs the information disclosure experiments. Although we have expected to see the opposite effect of previously demonstrated effects on privacy risk to affect on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure), neither the personality traits nor the preference scenario affected people's disclosure behavior. So the information disclosure 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.

Privacy vs. benefit. In this section, we investigated the influence of task design on actual information disclosure in the group explanation. We found that task design (i.e., the pressure on users to convince the group) affected participants' emotion-related information disclosure. This leads to the idea of privacy calculus (the tension between disclosure benefit versus privacy risk when disclosing information in a group explanation). Several studies show that when people want to decide on personal information disclosure, they trade off the anticipated benefits (i.e., the extent to which users believe disclosing their personal information to their group members is beneficial for the group decision or their negotiation position) with the risks of disclosure (i.e., the expectation of losses associated with the disclosure of personal information to the group), which is known as the "privacy calculus" [17, 70]. This remains for the next section to investigate this trade-off.

In the next section (see Section 5.2), we further unpack participants' actual disclosure decisions by measuring their perceived disclosure benefit versus privacy risk. Based on the literature, these factors might mediate the effect of the previously mentioned factors (i.e., personality, etc.) on participants' actual information disclosure.

5.2 Experiment 2: The Trade-off between Risk and Benefit on Information Disclosure

This thesis aims to study what makes good explanations for group recommendations. In Chapter 3, we started by determining what information needed to be conveyed to people in a group explanation for the recommended items. User comments highlighted that people might be sensitive to disclose some of their information in a group, e.g., their preferences. This led us to investigate the factors that we should model in the group to consider group members' privacy risk of information disclosure (see Chapter 4).

In an online experiment with real groups, we investigated the effects of three factors on people's privacy risk when disclosing personal information in the tourism context (see Section 4.2). We found that group members' personalities (using the 'Big Five' personality traits), their preference scenarios (i.e., having minority or majority preferences compared to two other group members), and the type of relationship they have in the group (i.e., loosely coupled heterogeneous, versus tightly coupled homogeneous) have a strong influence on people's perception of disclosure risk. Surprisingly, in contrast to this result, in information disclosure experiment (see Section 4.3), neither the personality traits nor the preference scenario affected people's actual disclosure behavior. However, upon further investigation in task design experiment, we found that the task design (whether group members were instructed to convince other group members of their opinion, or not) affected participants' emotion-related information disclosure.

So in this work, we try to investigate what other mediating factors could have been caused this gap between people's privacy attitudes compared to their disclosure behavior. For example could it be that perceived privacy risk (i.e., the expectation of losses associated with the disclosure of personal information to the group) and disclosure benefit (i.e., the extent to which users believe disclosing their personal information to their group members is beneficial for the group decision or their negotiation position) mediate the effect on participants' actual disclosure behavior. Several studies show that when people want to decide on personal information disclosure, they trade off the anticipated benefits with the risks of disclosure (i.e., [122]), which is known as the "privacy calculus" [17, 70]. Besides, in the group recommendation context, this effect depends on the task design (whether the task is designed as a competitive or cooperative task).¹ This thorough investigation of the dynamics between these factors and disclosure will result in a theory of user modeling that may inform considerations for generating group explanations automatically.

In this experiment, we find the intermediate factors that ultimately affect individuals' disclosure behavior from more general intermediate factors (i.e., individual's personality) to more specific intermediate factors (individual's privacy risk or disclosure benefit regarding one particular type of information). Results show that participants' personalities and whether their preferences align with the majority affect their perception of general privacy. This, in turn, affects their trust in the group, which affects their perception of privacy risk and disclosure benefit, which ultimately influence the amount of personal information they disclose in the group. We also find that *privacy risk* is a significant predictor when people are exposed to the competitive task but it is not for the cooperative

¹Note this is the same setting as in the task design experiment, but from user comments it is clear that this framing is more useful.

task.

- We design a conceptual model which helps us understand the relationship between privacy risk and disclosure benefit and their moderating factors (i.e., individual's personality) on actual information disclosure decision in the group.
- To create realistic scenarios of group decision-making where users can control the amount of information disclosed in the group and have iterative interaction between group members, we developed TouryBot. This web-based chatbot agent generates natural language explanations to help group members explain their suggestions to the group in the tourism domain.²

The contribution of this study is submitted to the User Modeling and User-Adapted Interaction (UMUAI) journal 2022 together with Geoff Musick, Bart Knijnenburg, and Nava Tintarev (under review).

5

5.2.1 Experimental Design

When people are in a situation where they have to decide where to visit next while traveling in a group, they must make a trade-off between disclosing their personal information to explain and support their arguments about what places to visit or to avoid (e.g., this place is too expensive for my budget) and protecting their privacy by not disclosing too much. This leads us to our main research question:

RQ How do people make a trade-off between disclosure benefit versus privacy risk in a tourism group recommendation scenario?

In this section, we describe an online, between-subjects study that investigates how antecedents of risk perception and perceived benefits relate to individuals' trade-off between **disclosure benefit** (i.e., disclosing their personal information to explain and support their arguments) versus **privacy risk** (i.e., not violating their privacy by revealing too much) when **disclosing their personal information** (e.g., their current location, emotion information, etc.) in a group recommendation explanation. Namely, we investigate: an individual's **personality**, **trust in group**, and **general privacy concern** as well as their **preference scenario**, and **task design**.³

Study Platform

To answer the research question, we implemented a web-based chat-bot that we call TouryBot. For the UI, we used a client in Java (Vaadin AI Chat)⁵ and implemented in the Vaadin

²More specifically, the initial POI options were selected from the category of "Food" in Amsterdam (see Section 5.2.1 for the details).

³All material for analyzing our results and replicating our user study, (i.e., user study materials, data gathered in the user study and the analysis scripts) is publicly available – (https://osf.io/z3hnp/?view_only=5db14a9c31ac4592bbdad98c5bbf7a3).

⁴The background color of the two chat windows (TouryBot chat vs. group chat) was selected to be different to help participants better to differentiate between the two chats.

⁵<https://github.com/alejandro-du/vaadin-ai-chat>, retrieved March 2021.



Figure 5.1: The chat in the peer majority and competitive task scenario between an active user and their group. Shown are the two UI sections: a) switching between two ongoing chats, one with a chat-bot and one with the group; and b) TouryBot suggests a place (the Oriental City in this example) for the whole group, one active user (John) and his two hypothetical group members (e.g., Bob and Alice) in a group chat.

5

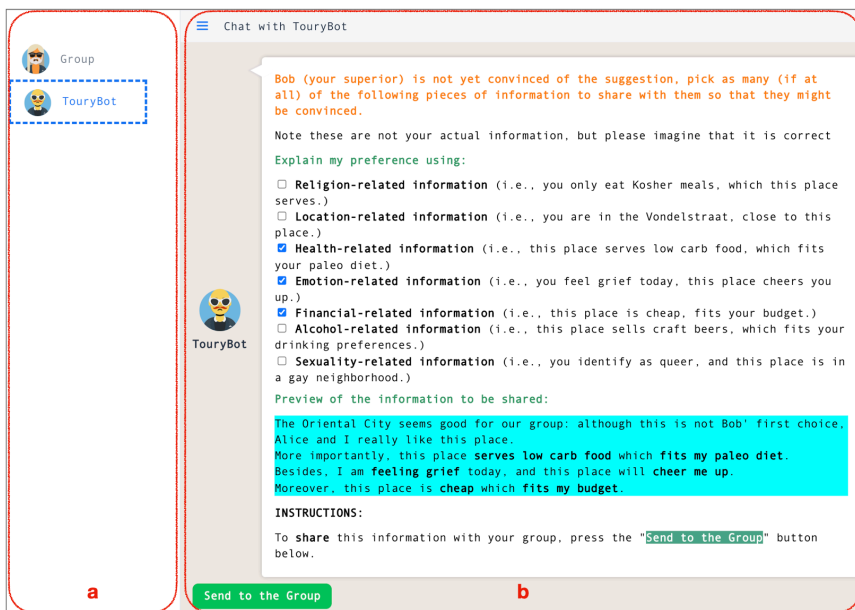


Figure 5.2: An example of chats where the active user is in the peer majority scenario (i.e., the user agrees with the majority preference with their peer) and is given a competitive task. Shown are the two UI sections: a) switching between two ongoing chats, one with a chat-bot and one with the group; and b) an *ongoing* chat with a chat-bot (TouryBot) where the user can indicate how much information they want to share to convince the other group member (Bob) to visit the suggested POI.⁴

framework.⁶ 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 with the Group (see Figure 5.1), and the other for the chat between the system bot and individual members (see Figure 5.2). Users can seamlessly switch between the two conversations to add system-generated recommendations and explanations to their discussions with other group members.

Manipulations

Inspired by previous experiments (see Sections 4.2 and 4.3), this experiment considers two factors that may influence information disclosure using between-subjects manipulations: users' preference scenario (3 conditions) and task design (2 conditions).

Preference scenario (binary). In a group size fixed to three, each participant in our study was exposed to either *minority* or one of two *majority* preference scenario types.

- **Minority:** 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 TouryBot.
- **Peer majority:** the active user's preference is in the majority within the group and against their superior. An item that is the user's favorite has been suggested to the group.
- **Boss majority:** the active user's preference is in the majority within the group and in line with their superior or boss. An item that is the user's favorite has been suggested to the group.

The shown scenario was dummy coded into two dichotomous values for both minority and peer majority tested against the boss majority.

Task design (binary). Each participant in our study was exposed to either a *competitive* or *cooperative* task design.⁷

- **Competitive task:** In this case, the participant tries to convince others to either skip or visit the recommended POI through privacy-sensitive explanations.
- **Cooperative task:** In this case, the participant is only tasked to reach a decision in their group through privacy-sensitive explanations.

Measures

Personal information disclosure, disclosure benefit, privacy risk, general privacy concern, trust in group, personality, and demographics were measured mainly using existing instruments. Except for demographic and personal information disclosure questions, all items were assessed using a 5-point Likert scale with endpoints of 'strongly disagree' and 'strongly agree'.

⁶An open platform for building web apps in Java (<https://vaadin.com/>), retrieved September 2021.

⁷Note this is the same setting as in the task design experiment, but from user comments it is clear that this framing is more useful.

Personal information disclosure. The primary dependent variable in our experiment is participants' personal information disclosure decision in a tourism group recommendation. To decide which personal information to include in the study, we used personal information categories listed in Caliskan et al. [9], which were derived from users' tweets on Twitter, and personal information used in Knijnenburg et al. [62], that used an online health application context. We included those that are relevant to a tourism recommender system context, namely the following personal information:

1. Emotion-related information (i.e., you feel grief today, this place cheers you up)
2. Location-related information (i.e., you are in the Vondelstraat, close to this place)
3. Financial-related information (i.e., this place is cheap, fits your budget)
4. Religion-related information (i.e., you only eat Kosher meals, which this place serves)
5. Health-related information (i.e., this place serves low carb food, which fits your Paleo diet)
6. Sexuality-related information (i.e., you identify as queer, and this place is in a gay neighborhood)
7. Alcohol-related information (i.e., this place sells craft beers, which fits your drinking preferences)

Users chose among these seven types of personal information as to which ones to share with their group members. In the final analyses, we consider all the information types as sum scores for the primary model analyses (the value ranges between 0 –when no information is disclosed at all–, to 7 –when all information is disclosed–). We consider disclosure as a sum score since the sharing selections are not independent decisions (i.e., when participants do the disclosure, they see all the information types simultaneously in a randomized order).

Disclosure benefit. In our context, perceived disclosure benefit refers to the “extent to which users believe disclosing their personal information to their group members is beneficial for the group decision”. To measure disclosure benefit for disclosing each type of information, we created seven questions, one for each of the seven personal information types that we included in the study as follows:

“I think disclosing my emotion-related information to these group members is beneficial for the group decision.”

The emotion-related information above was adapted based on the type of information asked. For the final analyses, the average disclosure benefit is centered on having a value between -2 to 2.

Privacy risk. Perceived privacy risk in our context is defined as the “expectation of losses associated with the disclosure of personal information in the group”. To measure privacy risk for disclosing each type of information, we created seven questions, one for each of the seven personal information types that we included in the study as follows:

Table 5.1: Items used to measure participants' privacy concerns and trust and the corresponding CFA outcome.

Factor	Item	Factor loading
General privacy concerns Alpha: 0.88 AVE: 0.550 Based on [63]	It usually bothers me when people ask me something personal.	0.742
	I will tell people anything they want to know about me.	0.678
	Compared to others, I am more sensitive about sharing personal information with other people.	0.726
	To me, it is the most important thing to keep things private from others.	0.754
	When people ask me something personal, I sometimes think twice before telling them.	0.699
	I think it is risky to tell people personal things about myself.	0.744
	I feel safe telling people personal things about me.	0.843
	I feel comfortable sharing my private thoughts and feelings with others.	0.734
Trust in group Alpha: 0.89 AVE: 0.761 Based on [50, 94, 121]	I trust the people in this group completely.	0.849
	I feel comfortable giving my personal information to the people in this group.	0.875
	The people in this group are trustworthy.	0.823
	The people in this group are honest.	0.904
	The people in this group are sincere.	0.908

"I think disclosing my emotion-related information to these group members is too sensitive for this type of group."

The emotion-related information above was adapted based on the type of information asked. For the final analyses, the average privacy risk is centered on having a value between -2 to 2 .

General privacy concern. This privacy concern is a personal trait pertaining to how concerned one is in general regarding their privacy. To measure general privacy concern we used the 8-item scale developed in Knijnenburg et al. [63] (listed in Table 5.1). Note this factor is scaled to have a variance of 1 in the final model.

Trust in group. By adopting the trust definition in Mayer et al. [78] to our context, the trust one individual has for another in the group can be defined as "the willingness of an individual to be vulnerable to the actions of other individuals by disclosing their personal information". To measure the active user's trust toward their group members, we adapted the items from previous research [50, 94, 121] as shown in Table 5.1. Note this factor is scaled to have a variance of 1 in the final model.

Factor loading of the included items for measuring *general privacy concern* and *trust in group* are shown in Table 5.1, as well as Cronbach's alpha and average variance extracted

(AVE) for each factor. The model has a good model fit: $\chi^2(64) = 271.991, p < 0.001$; root mean squared error of approximation (RMSEA) = .079; 90% CI : [0.070, 0.089], Comparative Fit Index (CFI) = 0.986, Tucker-Lewis Index (TLI) = 0.983. All included items have a higher factor loading than the recommended value of 0.40 [58]. Values for both Cronbach's alpha and AVE are good, indicating convergent validity, and the square root of the AVE is higher than the factor correlation, indicating discriminant validity of the two factors. The two factors are correlated with $r = -0.412$ (significant at $p < 0.001$).

Personality (continuous). We used the *Big Five Inventory* (BFI) to assess individuals' personality on the five traits of *Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism* [49]. 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.

Descriptive measures. We collected participants' age, self-identified gender, and nationality to enable a demographic description of our sample.

Materials

We needed some Initial Places Of Interest (POIs) in Amsterdam to elicit participants' preferences for the user study. To collect such POIs and make sure they somewhat fit participants' true preferences, we provided them with three initial POIs to rank. One POI among these three initial POIs was selected to recommend to the group in the TouryBot based on an active user's ranking on the three initial POIs. To encourage disclosure, we always recommended the active user's top choice in the majority scenario, and for the minority scenario, we recommended the user's least favorite place. The three initial POI options were retrieved from the most frequently visited POIs in the city of Amsterdam from the category of "Food", from the social location service *Foursquare*.⁸ Using participants' actual preferences, we aimed to increase the likelihood of a more realistic situation for users to imagine.

Procedure

Participants 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.

Step 1: "Group formation". Participants were asked for their first name and to form their (hypothetical) group by naming two people they might be in a group with whom they are not close. Further, participants were instructed to name members so that the social positions of the group members were not equal (e.g., a student planning a trip with a lecturer and another student or an employee planning a trip with a manager and another employee). This way, participants were in a hypothetical group with a "peer" and a "boss". Note the relationship type among group members was (in all cases) predefined as a "loosely coupled (weak ties) heterogeneous group" as described above (e.g., a lecturer

⁸<https://developer.foursquare.com/>, retrieved February 2021.

and students).⁹ Note that the group always consisted of three group members, where only one is the active user and two are *hypothetical* group members.

Step 2: “Preference Elicitation”. We also approximated active user preferences by asking them to rank three POIs in Amsterdam as described in Section 5.2.1.

Step 3: “Group Discussion”. Participants were randomly assigned to participate in one of our six scenarios (3 preference scenarios * 2 task designs).

Only one active user shares personal information to support their arguments in our setup. Depending on whether the current user is in the minority situation or one of the two majority type situations, they were tasked to convince other group members to skip or visit the suggested place in the competitive task design, or they were assigned to reach a decision in all three preference scenarios in the cooperative task design by disclosing personal information. As can be seen in Figure 5.1, in the *group chat* window, the recommendation came from the system (TourneyBot). This recommendation was based on the majority vote aggregation function in the group. After TouryBot suggested a restaurant for the group, the given active user was asked to switch to the *TourneyBot chat* window. As can be seen in Figure 5.2, in the *TourneyBot chat* window, the user was presented with different personal information options to support their arguments to the group (the background color of the two chat windows (group chat vs. TouryBot chat) was selected to be different to help participants better to differentiate between the two chats). They could choose which information (if any) that they wanted to share with their group to either persuade them or reach a decision with them. They could dynamically see the preview of the information to be shared with their group based on their choices. After they shared as much (if any) information as they wanted with their group members, the scenario ended with one of the hypothetical group members saying, “Okay, let’s skip/visit this place”. Then the participant was redirected to a questionnaire.

Step 4: “Questionnaire”. After completing the chat-bot activity, participants were asked a set of questions to assess their perceived general privacy concern, trust in the group members, personality traits, privacy risk, and disclosure benefit as described in Section 5.2.1.

Hypotheses

In Section 2.4, we discussed relevant literature on antecedents of privacy risk and disclosure benefit and ultimately their effect on actual personal information disclosure in group recommendation context. Based on that we formulate the hypotheses that lead this work as follows. Further, we develop a conceptual model to understand the relationship between those factors as can be seen in Figure 5.3. Inspired by single user recommender systems evaluation framework suggested by Knijnenburg et al. [60], we establish the core variables in the context of group decision making/recommendations as follows: *personal characteristics* of group members (i.e., their personality), *situational characteristics* with regards to the group (i.e., preference scenario), *subjective aspects* (i.e., their perceived risk when disclosing certain personal information in the group), and *actual behavior* of group members (i.e., when group members disclose their personal information in the group).

⁹Previously, in the *privacy factors* experiment (see Section 4.2), we found that privacy risk are perceived more in loosely-coupled heterogeneous groups than tightly-coupled homogeneous ones. In this work, we, therefore, focus on loosely coupled (weak ties) heterogeneous groups to consider privacy risk in an extreme case.




Figure 5.3: Conceptual model for antecedents of personal information disclosure.

- **H1)** Participants who perceive a higher level of disclosure benefits are more likely to disclose their personal information.
- **H2)** Participants who perceive a higher level of privacy risk are less likely to disclose their personal information.
- **H3)** Participants who perceive a higher level of trust in the group members perceive higher levels of disclosure benefit.
- **H4)** Participants who have been told that the group decision is a competitive task perceive a higher level of disclosure benefit than participants who have been told to address the decision as a cooperative task.
- **H5)** Participants whose preferences are in the minority perceive higher disclosure benefits compared to participants in both majority scenarios.
- **H6)** Participants who have a higher level of trust in the other group members perceive lower levels of perceived privacy risk.
- **H7)** Participants whose preferences are in the minority perceive higher levels of privacy risk compared to participants in both majority scenarios.
- **H8)** Participants with a higher level of general privacy concern have less trust in their group members.
- **H9)** Extraversion affects participants' general privacy concern perception.
- **H10)** Agreeableness affects participants' general privacy concern perception.

- **H11)** Consciousness affects participants' general privacy concern perception.
- **H12)** Neuroticism affects participants' general privacy concern perception.
- **H13)** Openness affects participants' general privacy concern perception.

5.2.2 Results

In this section, we discuss the outcomes of the hypothesis tests we conducted and present exploratory findings. We built a structural equation model (SEM) upon the data collected with our questionnaire by using the R library Lavaan¹⁰.

Participants

To determine the required sample size, we performed a power analysis [19] of a small-sized effect (0.2 SD) with a power of 85% in a between-subjects experiment. It showed that a minimum of 277 participants were needed in total. This was in line with the suggested minimum sample size for SEM in Knijnenburg et al. [58] (minimum 200 participants).

We recruited 280 participants from the crowdsourcing platform *Prolific*.¹¹ This platform has shown to be an effective and reliable choice for run (two participants were excluded from the results). The resulting sample of 278 participants had an average age of 25.8 (sd = 7.5) with a satisfactorily balanced gender distribution (49% female, 50% male, and 1% other).

Hypothesis Tests

The resulting SEM model (Table 5.2) shows how *privacy risk* and *disclosure behavior* and their antecedents influence personal information disclosure in groups. Based on the final results, all the question items to measure general privacy concern and trust in group (see Section 5.2.1) remained valid. The model has a great model fit: chi-square(204) = 364.150, $p < 0.001$; root mean squared error of approximation (RMSEA) = .053; 90% CI : [0.044, 0.062], Comparative Fit Index (CFI) = .987, Tucker-Lewis Index (TLI) = .992. We looked at the disclosure and respective privacy risk and disclosure benefit as an average rather than on an individual item level because when people do the disclosure, they see all the information types simultaneously and probably reason about what they will disclose (not independent decisions).

The results show that the relationship between *average disclosure benefit* and *overall disclosure* is significant ($\beta = 0.927$, $p < 0.001$), supporting H1. Furthermore, a significant negative interaction effect of *average privacy risk and task design* on *overall disclosure* can be observed ($\beta = -0.420$, $p = 0.043$). This finding suggests that when average disclosure benefit seems to be the same (as there is a weak but strongly significant correlation between the average disclosure benefit and average privacy risk, $\beta = -0.080$, $p < 0.001$), average privacy risk has a significantly stronger effect when it is a competitive task versus a cooperative task on the overall disclosure. Given that we have a significant interaction effect, the main effects cannot be interpreted in isolation. Therefore H2 is supported, but with the caveat that it depends on task design (Section 5.2.2 describes the interaction effect of risk and task design on disclosure in more detail).

¹⁰<http://lavaan.ugent.be/>, December 2021

¹¹<https://www.prolific.co>

Table 5.2: The results of the SEM analysis of the final model. We tested these factors as part of our hypothesis tests (see Section 5.2.1).

Hypothesis	Standardized Estimates	Standard Error	P-Value	Supported?
H1: Benefit ->Disclosure (+)	0.927	0.137	0.000	Yes
H2: Risk ->Disclosure (-)	-0.343	0.111	0.002	Yes
H3: Trust in Group ->Benefit (+)	0.178	0.026	0.000	Yes
H4: Task Design ->Benefit	0.035	0.060	0.566	No
H5: Minority ->Benefit	0.102	0.076	0.161	No
H5: Peer Majority ->Benefit	-0.114	0.076	0.118	No
H6: Trust in Group ->Risk (-)	-0.123	0.031	0.000	Yes
H7: Minority ->Risk	0.079	0.086	0.341	No
H7: Peer Majority ->Risk	-0.026	0.086	0.750	No
H8: General Privacy Concern ->Trust in Group	-0.358	0.058	0.000	Yes
H9: Extroversion -> General Privacy Concern	-0.536	0.092	0.000	Yes
H10: Agreeableness -> General Privacy Concern	-0.629	0.127	0.000	Yes
H11: Conscientiousness -> General Privacy Concern	0.550	0.130	0.000	Yes
H12: Neuroticism -> General Privacy Concern	0.093	0.090	0.301	No
H13: Openness -> General Privacy Concern	-0.154	0.140	0.272	No

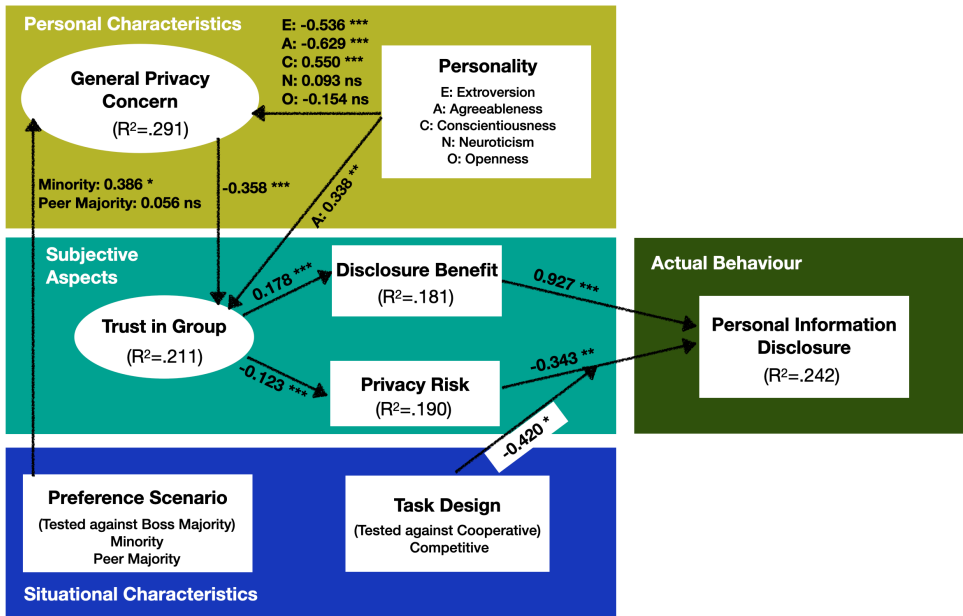


Figure 5.4: The Structural Equation Model (SEM) for the data of the experiment. The model shows the subjective factors behind users' information disclosure decisions when using a group recommender system, and the effect of personal and situational characteristics (Significance levels: *** $p < .001$, ** $p < .01$, * $p < .05$, 'ns' $p > .05$).

Besides, *task design* is not a significant predictor of *average disclosure benefit* ($\beta = 0.035$, $p = 0.566$), and therefore, H4 is not supported. Considering *preference scenarios*, it is not a significant predictor of *average disclosure benefit* ($\beta = 0.102$, $p = 0.161$) or *average privacy risk* ($\beta = 0.079$, $p = 0.341$), and H5 and H7 are not supported respectively. Furthermore, the analysis results show that *trust in group* has a significant impact on both *average disclosure benefit* and *average privacy risk*, supporting H3 ($\beta = 0.178$, $p < 0.000$), and H6 ($\beta = -0.123$, $p < 0.000$). Moreover, *general privacy concern* negatively affects *trust in group* ($\beta = -0.358$, $p < 0.000$), supporting H8. We found that high levels of the *agreeableness* trait also has a significant positive effect on *trust in group* ($\beta = 0.338$, $p = 0.004$).

The analysis results indicate that three out of five types of *personality traits* (*extroversion*, *agreeableness*, and *conscientiousness*) are related to *general privacy concern* (the first two negatively and the last positively), supporting H9 ($\beta = -0.536$, $p < 0.001$), H10 ($\beta = -0.629$, $p < 0.001$), and H11 ($\beta = 0.550$, $p < 0.001$). However, the relationship between the other two *personality traits* (*neuroticism*, and *openness*) and *general privacy concern* is not significant, and so H12 ($\beta = 0.093$, $p = 0.301$) and H13 ($\beta = -0.154$, $p = 0.272$) are not supported. Additionally, *minority preference scenario* is found to have a significant positive relationship with *general privacy concern* ($\beta = 0.386$, $p = 0.018$). Figure 5.4 summarizes the final model.

We commit to make all data and code publicly available for the community to be able to replicate and reproduce our study and results.¹² However, the raw results of our user

¹²<https://osf.io/z3hnp>




Figure 5.5: The jittered scatterplot displaying the distribution of perceived average privacy risk (x-axis) and the corresponding overall information disclosure (y-axis) in two different tasks (competitive vs cooperative), with estimated regression lines.

5

Table 5.3: Participants’ level of disclosure, average disclosure benefit, and average privacy risk per item.

Items	Emotion	Location	Financial	Religion	Health	Sexuality	Alcohol
Level of disclosure	65(23%)	209(75%)	172 (62%)	25(9%)	116(42%)	28(10%)	78(28%)
Disclosure Benefit	3.3	4.0	3.6	2.7	3.8	2.4	3.0
Privacy Risk	3.1	2.0	2.7	2.9	2.4	3.1	2.6

study are anonymized, i.e., we do not publish participants’ identifiable information such as user IDs.

Exploratory Findings

Here, we present several exploratory findings that may help explain the results of the hypothesis tests.

Interaction effect of risk and task design on disclosure

Figure 5.5 visualizes the distribution of perceived average privacy risk (x-axis) and the corresponding overall information disclosure (y-axis) in two different tasks (competitive vs. cooperative), with estimated regression lines. For example, the red line shows as risk increases users disclose less information in the competitive task. As seen in the figure, in line with our findings, there is a negative slope for the cooperative task between overall disclosure and perceived privacy risk, but this slope is smaller than it is for the competitive task.

Disclosure behaviour per item

Here we look at each item individually to see how much of each type of information that participants disclosed and how much benefit and risk they perceived regarding it. As can

be seen in Table 5.3, people disclosed location and financial the most (75% and 62% respectively) and sexuality and religion the least (10% and 9%). Among all seven information types, it seems people found the location, health, financial, emotional, and alcohol information more beneficial for this context to share with group members (disclosure benefit ≥ 3). They perceive more privacy risks when disclosing emotional and sexual information.

Qualitative feedback

Participants distinguished between the cooperative and competitive tasks in their qualitative comments. For example, in the cooperative task, when they were asked why they disclosed certain information, some of them explicitly mentioned a group goal as shown below:

“Because it can help in making an efficient decision.”, “Make them aware of my whereabouts, so they come to make the right decision.”, “It is practical information that might help in choosing the right attraction sites for the group.”, or “If you know/trust the group, you can help them decide where to go.”. “Nowadays more and more people start to live a healthier life so disclosing health-related information could be very useful to make a decision.”.

Comparatively, in the competitive task, people seemed to follow a more self-serving (egocentric) goal to either not disclose or disclose certain information. Some examples include the following: *“I don’t think it’s fair to persuade someone to go somewhere based on my location as opposed to theirs. But if we all had the same location roughly, that would be fine.”, “Everyone is different, and what might be classed as the perfect diet to one person may be viewed as boring and restrictive to others, therefore I didn’t feel this was a valid argument in this case as I didn’t know the people.”, “I think they wouldn’t care and would just think I’m too picky.”, “Health is extremely important, and I would not be willing to put myself in a situation where something will compromise my health.”, or “As for the religion subject, even if it’s personal, it has a big interest since it could stop me from eating or could lead to me getting sick.”.*

User comments indicate that they saw the task as competitive or cooperative and that this informed the reasoning behind the disclosure. This can inform design regarding formulating tasks in group recommender systems (e.g., focus on consensus and cooperation when asking people for personal information).

5.2.3 Discussion

The study results provide exciting insights into users’ personal information disclosure decisions in a tourism group decisions/recommendations context. They also demonstrate how personal privacy risk perception and manipulated situational characteristics influence the decision process. In this section, we reflect on these results, their design implications, and the limitations of our study.

Establishing trust in the group is essential. Disclosure is a trade-off between risk and benefit that is rooted in trust. When people have to decide where to visit next while traveling in a group, the decision-making facilitator like the one we proposed performs better if trust in the group is high. In that case, people perceive less privacy risk and more disclosure benefit and ultimately disclose more personal information to help group decision-making. It has also been shown that higher degrees of trust at the individual and

group levels help group members implement more effective and meaningful processes to make collective decisions [113]. Thus, we recommend making sure to establish trust within the group beforehand. There are suggestions in different domains to facilitate trust and active participation among group members by, for example, taking their opinions on the decisions into account [67].

Interaction effect of risk and task design on disclosure. Although we expected to see an effect of task design on disclosure benefit, we only found a significant negative interaction effect of average privacy risk and task design on overall disclosure. It could be that for the selected task design in this study (or how the two task types were operationalized), users did not perceive any distinguishable difference in the benefit of disclosing their information. Future work should investigate this with different types of task design or even the same ones with different formulations. Regardless, the current study shows that how the group decision task is framed (i.e., cooperative or competitive) can have a substantial impact on *how* people make privacy decisions—in particular, it influences the importance of *risk* in the decision. While the average level of privacy risk was roughly the same (2.7 out of 5) between participants in the competitive and cooperative task conditions, there is a significant interaction effect of task design and privacy risk on information disclosure: privacy risk has a significantly stronger effect on the disclosure decision when it is framed as a competitive task (as compared to a cooperative task).

A potential explanation for this effect could be that the cooperative task was viewed as having a more altruistic goal, while the competitive task was seen as having a more self-serving goal, as can be seen through participants' qualitative feedback in Section 5.2.2. In the competitive task, the information disclosure is thus for one's own benefit, hence people will weigh their personal risk regarding the information disclosure with how much benefit they think they are going to get out of this disclosure. Toma and Butera [129] stated that competition activates the fear of being exploited (risk vulnerability), but also the desire to exploit other people. They also add, in all information exchange situations, competition activates tactical deception tendencies aimed at maintaining a positive self in other people's eyes [129]. In contrast, when people regard the information disclosure as benefiting the group (i.e., in the cooperative task), then it seems one's own privacy risk becomes a less critical factor which leads to more disclosure—one that can be sacrificed for the good of the group.¹³ In line with this, user comments give us an idea of how to formulate tasks in group recommender systems. As such, when designing for situations where disclosure is crucial for the success of a system, designers should emphasize cooperative aspects of the system's goal in their communication to the users as in this case people are more likely to disclose personal information (e.g. "Help make the recommendations better by providing some information about you / your preferences").

Effect of preference scenario on general privacy concern. Although we expected that the preference scenario would have a direct effect on privacy risk and disclosure benefit, there are two other mediating factors in between (i.e., first general privacy concern and

¹³This suggests that people would expect others to reciprocate this behavior. A future study with repeated opportunities for mutual disclosure could investigate whether this influences participants' behavior in the long run.

then trust). It is counter-intuitive that group members' general privacy concerns (which are often considered to be a stable personal trait) could have been influenced by our manipulation of the preference scenario (i.e., whether their preferences are aligned or not aligned with the preferences of the majority in the group). However, as we measured privacy concern right after the experiment, the presented scenario might have had a lingering effect on participants' expression of that concern, even though the questions were asked more generically. In particular, our study finds that *when people are in a scenario where their preferences do not reflect those of the majority, they perceive significantly higher privacy concerns compared to people whose preferences are aligned with the majority* (regardless of whether this means that they are siding with a peer or with a superior). Their increased concerns, in turn, have a negative effect on their trust in the group, which influences their perception of risk and benefit, which may ultimately reduce the amount of information they disclose.

5

Effect of personality traits on privacy. Our results indicate that extraverts have lower privacy concerns; people with high agreeableness have lower privacy concerns and higher trust; conscientious people have higher privacy concerns; however, there's no effect of neuroticism and openness. The findings of the effects of extraversion [104], agreeableness [1, 104], and conscientiousness [104] are aligned with previous works, while other findings regarding neuroticism and openness are not. Page et al. [100] give a potential explanation for the inconsistent effects of personality on privacy concerns: in most research personality serves as a crude proxy for more specific personal characteristics—such as “communication styles”—that have a much closer relationship with privacy concerns. Using more specific personal characteristics remains open for future work.

Effect of agreeableness on trust. Our results indicate that high levels of the *agreeableness* trait has a positive effect on *trust in group* and a negative effect on *general privacy concern*. Agreeableness “involves getting along with others in pleasant, satisfying relationships” [101]. Agreeableness emphasizes trust, altruism, compliance and modesty [1]. Agreeable individuals are also less likely to judge others' actions as potentially harmful when faced with privacy threats. Hence, their tendency to trust and to be less suspicious of their environment may reduce their privacy concern. Consequently, they may have lower privacy concerns [104].

We conducted two experiments discussed in this chapter looking at the factors that influence individuals' disclosure behavior in the group to answer how people trade-off between disclosure benefit versus privacy risk to decide on information disclosure in a group explanation. The following section (see Section 5.3) summarizes our findings of the experiments in this chapter.

5.3 Chapter Conclusions

In a group recommendation/decision context, there is information that we can present to people in a group to help them reach a consensus on the recommended items. Adding more detailed information in situations when there is disagreement in the group seems more beneficial. However, this increases individuals' privacy risk perception (it is measured by a possible loss of privacy as a result of the system presenting an explanation including people's personal information to the group) and ultimately decreases their information disclosure in the group (it is measured by which personal information, if any, users share with their group members). This chapter covered two experiments looking at the factors that influence individuals' disclosure behavior in the group.

Experiment 1: task design. In the privacy factors experiment (see Section 4.2), we found three factors that influence the group members' perceived privacy risk in the tourism group recommendation context, namely, group members' *personality*, *preference scenario*, and the type of *relationship*. In a follow-up experiment (the information disclosure experiment, see Section 4.3), although we have expected to see the opposite effect of previously demonstrated effects on privacy risk to affect on information disclosure (i.e., if a factor increases user privacy risk, it decreases their information disclosure), neither the personality traits nor the preference scenario affected people's disclosure behavior. The information disclosure experiment required participants to *convince* other group members. Therefore, one explanation for why we did not obtain the expected results could have been that our task design nudged participants into a "convincing mindset". So in this experiment (see Section 5.1), we investigated the effect of task design on people's disclosure behavior.

We found that the task design (whether group members were instructed to convince other group members of their opinion or not) affected participants' emotion-related information disclosure. There might be additional mediating factors that influence participants' actual disclosure behavior. In the following experiment, we investigated other intermediate factors (i.e., individual's personality as well as their preference scenario and the task design) together, besides perceived privacy risk and disclosure benefit as these additional measures might mediate the effect on participants' actual disclosure behavior.

Experiment 2: benefit vs. risk. Our contradictory results between the privacy factors and the information disclosure experiments (where factors, i.e., individual's personality as well as their preference scenario that influenced privacy risk did not consistently affect disclosure behavior, see Section 4.2 and Section 4.3) could also be explained by a tension between benefits and risks. In fact, several studies show that when people decide on personal information disclosure, they trade off the anticipated benefits with the risks of disclosure. This leads us to the *benefit vs. risk* experiment (see Section 5.2). We presented an online user study investigating how individuals trade-off between *disclosure benefit* versus *privacy risk* to decide on *information disclosure* in a group recommendation explanation. For example, how individuals trade-off between disclosing their personal information to explain and support their arguments while not violating their privacy by revealing too much information regarding their current location, emotion, etc., in a group. We specifically study how the following antecedents of risks and benefits relate to this

trade-off: an individual's *personality* (using the 'Big Five' personality traits), *trust in group* (i.e., whether group members trust the other individuals in the group), *general privacy concern* (i.e., an individual's general tendency to worry about information privacy) as well as their *preference scenario* (i.e., having minority or majority preferences compared to two other group members), and *task design* (i.e., either instructed to convince the group to visit or skip a recommended place or, to reach a decision in the group). Note that when designing the experiment, we thought of task design in terms of convincing people as in the task design experiment. However, in the new study design, we saw in user comments that they perceived that task more as cooperative and competitive settings (see Section 5.2).

This experiment helps us formulate a model that shows which intermediate factors affect individuals' disclosure behavior, from more general intermediate factors (i.e., individual's personality) to more specific intermediate factors (individual's privacy risk or disclosure benefit regarding each particular type of information). Results show that preference scenarios and one's personality (extroversion, agreeableness, and conscientiousness) change their general privacy concern. The general privacy concern in turn significantly affects their trust perception of the other group members. The trust in turn changes their perception of privacy risk and disclosure benefit in two different directions, slightly higher effect in benefit than risk. Ultimately there is a negative interaction effect of average privacy risk and task design, and a positive effect of benefit on people's information disclosure behavior in the tourism group recommendation context.

5

Wrap up

- Due to the diverse needs and preferences, recommendations for *groups* are particularly challenging and often require discussions among group members. So to create realistic scenarios of group decision-making where users can control the amount of information disclosed, I developed and provided an open-source web-based chatbot, TouryBot.
- To protect group members from unwanted personal information disclosure in their group, we should avoid urging people into a "convincing" mindset.
- We propose a conceptual model for transparency which considers different factors. Through this model, we develop a better understanding between perceived risk and benefit and how they interact with disclosure. To model this, we conduct a comprehensive statistical evaluation of the results using structural equation modeling (rather than a series of regressions).
- We found privacy risk is essential and different between competitive and cooperative tasks. This result suggests that, We should frame the task of finding a suitable destination in the group cooperatively when asking people for personal information; in this case, people are less likely to perceive risk. The literature states that competition activates the fear of being exploited (risk vulnerability) and the desire to use others. When they are doing it for the group (i.e., in the cooperative task), then it seems their own privacy risk becomes less critical.
- We should establish trust in the group beforehand (i.e., that group members trust the other individuals in the group) because the decision-making facilitator like the one

we proposed performs better if trust in the group is high. When trust is high in the group, people perceive less privacy risk and more disclosure benefit and ultimately disclose more personal information, which helps groups make informed decisions. This can be reached, for example, by taking the opinions of every group member on the decisions.

This thesis represents a step towards developing explanations for group recommendation/decision systems by taking group members' privacy concerns into consideration. The next chapter (see Chapter 6) summarizes our findings in this thesis and provides design recommendations for future work for designing explanations for group recommendations.

6

Conclusions and Future Work

In this thesis, we have explored what makes good explanations for group recommendations by considering group members' privacy concerns. We mainly focused on studying the impact of different individual and situational characteristics on people's privacy risk perception regarding group explanations and modeling people's actual disclosure decisions in this context. There are both benefits for disclosing information in a group explanation as well as risks. Some studies that only study information disclosure may not be able to detect this tension. Based on a more profound understanding of users' disclosure behavior, the core contribution of this thesis is a privacy disclosure model. This model contains different individual and situational models/characteristics that help predict users' disclosure intention in a group decision/recommendation context. In this last chapter, we summarize the main findings of what needs to be considered when generating explanations for groups, reflect on the work carried out in this thesis, and outline future research directions.

6.1 Research Questions Revisited

In this section, we revisit and answer the three sub-research questions introduced in Chapter 1 to answer our main research question: *How do different factors influence individual group members' requirements towards a group explanation?*

Note that the following sub-research questions have evolved in the process of my Ph.D. project.

RQ1 What information should be disclosed in a group explanation to increase group members' satisfaction?

To answer **RQ1**, which addresses what information should be conveyed when generating a group explanation, we started with assessing different aggregation strategies used to recommend items to groups (see Section 3.1). As with explanations for individual recommendations, explanations for groups can be designed based on the underlying recommendation algorithm. The aggregation strategies (aggregate individual item-ratings predictions which are given) are a form of recommendation algorithm for groups, which is our basis in this thesis for the explanations that should be developed for group recommendations. In the aggregation strategies experiment (see Section 3.1), we presented a user evaluation of different explainable aggregation strategies. Based on the results, it seems that it does not matter which aggregation strategy we use in terms of user satisfaction as long as we explain it. So one way to understand whether the explanation is necessary in the first place is to separate the output of the aggregation strategy from the output of the explanation, which was investigated in the next experiment. So the aggregation strategies versus explanations experiment (see Section 3.2) evaluates users' perceptions regarding aggregation strategies and their explanations separately (in isolation). It seems that explanations containing information about the aggregation strategy do not significantly benefit group members (i.e., increase group members' satisfaction) in simple scenarios like the one we used (i.e., a few candidate items to choose from, and a smaller group of users). However, these experimental results are not enough to claim that explanations, in general, are not helpful for group recommender systems. More complex scenarios might involve a more balanced situation between subgroups of people/users with different preferences or a greater number of options to choose from; or group members have item rating disagreement in such cases, an explanation of the approach used might have an impact. Therefore, the formulating group explanations experiment (see Section 3.3) proposes different group explanation styles for more complex scenarios, for example, when a group member did not receive her favorite item and for a great number of candidate items (10 items compared to 3 items in the aggregation strategies versus explanations experiment). In that study, user comments highlight the need to protect certain types of information, i.e., group members' ratings of items. This suggests that studying the other factors, such as privacy (i.e., protecting certain types of information from the group) appears to be a more promising direction than explaining aggregation strategies. This motivates our second research question **RQ2**.

RQ2 How do different factors (i.e., individual differences, group dynamics, etc.) influence individual group members' privacy risk perception of information disclosure in a group explanation?

Grounded on the findings of Chapter 3, to answer **RQ2** which investigates the factors that we should model in the group to consider group members' privacy, we decided to go deeper into what people would disclose in a group explanation. In the privacy preferences experiment (see Section 4.1), we find that there may be some individual differences in the level of privacy risk. This leads us to the privacy factors experiment (see Section 4.2), where we investigate which factors contribute to these individual differences. We studied some factors identified in the literature that influence individual privacy risk. Our results confirmed that these factors seem to matter in our experiment's context as well, namely: **a**) group members' personality (using the 'Big Five' personality traits), **b**) specific preference scenarios (i.e., whether the user's preferences are aligned or not aligned with the preferences of the majority in the group), **c**) the type of relationship they have in the group (both relationship strength and equality of positions, i.e., loosely coupled heterogeneous, versus tightly coupled homogeneous). The effects of these factors on group members' privacy risk perception were measured regarding all of the selected information types (e.g., participant's current location, emotion, financial, religion, health, sexuality, and alcohol-related information) included in a single group explanation. Still, the participants' comments suggested individual differences regarding which information to disclose in relation to the three factors. So we decided to study the influence of the user model we predicted from the privacy factors experiment on disclosure for two specific information types (location and emotion)¹. We studied disclosure for the two information types separately, not in a single explanation as in the information disclosure experiment (see Section 4.3). I developed a web-based chat-bot called TouryBot for this experiment to create realistic scenarios of group decision-making where users can control the amount of information disclosed. This chat-bot agent generates natural language explanations² to help group members explain their arguments for or against the places suggested to the group. Surprisingly, we did not find an effect of personality, preference scenario, or relationship type on group members' disclosure decisions. This result can be explained by the study design, as our study asked users to convince other group members to accept or skip the suggested place. Therefore, one explanation for why we did not obtain the expected results is that our task design nudged participants into a "convincing mindset". The convincing mindset might cause participants to disclose more than they would want, and perceived risk does not get expressed. Therefore, in Chapter 5, we study the effect of task design and other potential intermediate factors on disclosure behavior.

¹We selected location and emotion information types among the five types of information we included in the explanation in the privacy factors experiment (see Section 4.2). These two are the most used information types in current tourism recommender systems (e.g., [82]).

²In this thesis, I use a template-based *natural language generation* technique by adding pre-defined templates (i.e., contains a controlled vocabulary) which can be easily adapted and extended based on the user's selected options.

RQ3 How do people trade-off between disclosure benefit versus privacy risk of information disclosure in a group explanation?

Based on what we discussed in the information disclosure experiment, to answer **RQ3**, we decided to go deeper into *task design* (see Section 5.1). The information disclosure experiment asked users to do 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. So the task design (convincing mindset) could have decreased the effects of factors on information disclosure – which have been demonstrated in the privacy factors experiment – to the degree that we could not pick them up in the information disclosure experiment. Therefore, in the task design experiment, we studied the effects of task design, using the TouryBot introduced before. We found that task design affected participants' disclosure decisions (for certain types of sensitive information). So based on the task design experiment's results, we realized that the effect of task design on disclosure behavior is significant. This leads us to investigate the intermediate step using the TouryBot in terms of privacy risk and disclosure benefit and their antecedents factors (i.e., personality, etc.), not just in terms of what people disclose in the benefit vs. risk experiment (see Section 5.2). There might be internal processes that happen that we do not necessarily see in user disclosure behavior. This experiment helped us formulate a privacy disclosure model that shows what factors affect individuals' disclosure behavior, from more general intermediate factors (i.e., individual's personality) to more specific intermediate factors (individual's privacy risk or disclosure benefit regarding one particular type of information). Results indicate that individuals' personality (using the 'Big Five' personality traits) and their preference scenario (whether one's preferences align with the majority in the group) affect their perception of general privacy. This, in turn, affects their trust in the group, which affects their perception of privacy risk and disclosure benefit when disclosing personal information, ultimately influencing the amount of personal information they disclose. Besides, an interesting finding indicates that privacy risk on information disclosure is different for different types of tasks, specifically for a competitive task (i.e., instructed to convince the group to visit or skip a recommended place) versus a cooperative task (i.e., instructed to reach a decision in the group).³ Privacy risk significantly impacts information disclosure when the task of finding a suitable destination is framed competitively but not when it is framed cooperatively. This suggests the cooperative formulation of tasks can be used in group recommendations to facilitate reaching consensus.

6

6.2 Practical Considerations

In this section, we outline how scientists and developers could use the work in this thesis to further the development of privacy-preserving explanations for group recommendations.

³Note that when designing the experiment, we thought of task design in terms of convincing people as in the task design experiment. However, as we saw in Section 5.2, the participants' perception showed in the new study design is more about cooperative versus competitive settings.

6.2.1 Empirical Implications

- We recommend adapting the group explanation style (i.e., repairing vs. reassuring) to the variation in user preferences (i.e., when there is group disagreement or agreement on the recommended item). When all group members agree on the recommended item, we recommend using reassuring explanations and keeping it short. Otherwise, when people have different preferences in the group, using repairing style explanations that focus on why the item is recommended and including more details seems more beneficial (see Chapter 3).
- When the recommended item does not align with the majority preferences, we recommend being cautious about disclosing the identity (i.e., name) of the people with the minority preferences in the group, together with their strong opinions, but instead making it anonymous. For example, *some of the group members want to visit this place and won't be talked out of it easily* (see Chapter 4).
- We recommend making sure to establish trust (i.e., whether group members trust the other individuals in the group) within the group for group decision-making. When trust is high in the group, people perceive less privacy risk and more disclosure benefits and ultimately disclose more personal information to help group decision-making (see Chapter 5).
- The task of finding a suitable destination in the group should be framed cooperatively when asking people for personal information. As in this case, people are less likely to perceive the privacy risk of personal information disclosure, which leads to more disclosure. The literature states that competition activates the fear of being exploited (risk vulnerability) and the desire to use other people. When they are doing it for the group (i.e., in the cooperative task), then it seems their own privacy risk becomes less critical, and they might be willing to disclose more personal information to the group members (see Chapter 5).
- To develop a good explanation for group recommendations in terms of our goals, we should first understand the group and its dynamics well. One-size-fits-all explanation formulation should not be used for the group decisions/recommendation context as our results show privacy concerns and disclosure behavior in groups is very context-dependent and depends on the individual group members' characteristics as well as situational characteristics. We recommend, among other potential factors, considering group members' personalities, whether their preferences align with the majority, their social relationship in the group, their trust in the receiver (or, in our group recommender context, the group), task design (whether the task of finding a suitable item for the group is framed cooperatively or competitively), their general privacy concern perception, their perception of privacy risk and disclosure benefit to ultimately predict the amount of personal information they would disclose in such a situation (see Chapter 5).

6.2.2 Methodological Lessons Learned

- For experiments aimed at benchmarking, we recommend a thorough reporting on how participants were recruited in a study as any selection of study participants can influence the evaluation outcome, which should not be generalized outside the scope of the scenario (see Chapter 3).
- We recommend ensuring consistency in measurement or describing and motivating changes well to benchmark user studies (see Chapter 3).
- We recommend that future work not only describes the cases where explanations can be generated but also represents the edge cases for which they cannot (see Chapter 3).
- The recruitment of group participants for group experiments is a big challenge, especially when aiming to control the group type. The challenge increases when recruiting heterogeneous, loosely coupled groups (i.e., a group of colleagues including a manager). A recommendation for future studies is to first ask the participants from a “higher” position rather than recruiting organically or requesting participants in “lower” positions to form the group. We received feedback from several participants that it is difficult for them to ask a person in a higher position (e.g., their boss) to form a group with them.

6

6.3 Limitations and Future Directions

Designing explanations for group decisions/recommendations still has room for improvement. Here, we discuss the limitations of our experiments and highlight a few promising directions that can advance the design of explanations in the group decisions/recommendations context.

Hypothetical personal information. We measured participants’ privacy risk, disclosure benefit, and actual disclosure behavior regarding hypothetical personal information rather than their actual personal information (e.g., their current location, emotion, financial, religion, health, sexuality, and alcohol-related information). The use of hypothetical information allowed us to avoid privacy concerns with the experiment itself (which could have resulted in a participant selection bias) and the effect of individual differences in the sensitivity of participants’ actual personal information (e.g., someone with an alcohol addiction will likely find their alcohol-related information more sensitive than someone who does not drink alcohol). A downside of using hypothetical information is that our participants may have been unable to imagine the situation or that they behaved differently from how they would have behaved if the disclosure scenario presented in the experiments considered their actual profile. Although we asked them to imagine that the experiments scenario considered their real information, and participants’ answers to the open-ended questions show their high engagement in the experiments, asking participants to share hypothetical personal information still might lead to different results than if the experiments would consider their actual personal information. Future work could attempt to replicate our findings in real-world group decision-making settings.

Hypothetical group. A related limitation is that our scenario involved hypothetical group members (except for one experiment involving users in real groups). Because asking people to form a group was challenging as not every member would accept to join or continue until the end of the experiment. So even if one group member dropped out, we needed to discard data for the whole group. To reduce the recruiting cost and the complexity of the experiments, each group contained only one active user (participant), who was asked to imagine a specific group based on the specified criteria. To increase the realism of the scenario, we asked participants to enter the actual names of the people they imagined to be in this group. We used those particular names throughout the experiment. Future work could study people in real groups to see how group members with different disclosure behaviors interact in a privacy-preserving way to reach a consensus.

Group size. To simplify the design of our experiments, the presented scenario always involved a group of exactly three people (not an uncommon group size). Future work could investigate how group size affects the outcomes of our experiments. The effect of group size is not trivial. For example, a larger group means that any disclosure reveals data to more people, which may increase the potential privacy risk. On the other hand, a larger group also means that more people disclose their personal information hence one's own information may be sheltered in the sea of information. Larger groups also have the potential to result in information overload. In that situation, recommending what information is more important to justify one's opinion becomes more critical since giving too much information in the justification might cause it to be ignored.

Recommendation domain. The actual disclosure decision might be domain-dependent. For example, low involvement and high involvement decision domains [105]) could be perceived differently in terms of privacy risk and disclosure benefit. Depending on the type of item, users tend to invest more or less time until a final decision is taken; items with high related decision efforts are marked as high-involvement items, whereas items with low related decision efforts are marked as low-involvement items. For example, in a high-involvement decision domain like the choice of a shared apartment, people might perceive that disclosing personal information has more benefits if it helps to make a better group decision. At the same time, high-involvement domains may require disclosing more sensitive personal information (e.g., in the case of a shared apartment, budget information). The current study was conducted in the context of tourism—a domain suitable for studying group decisions/recommendations, as it is relatable for many participants and commonly involves coordinating with a group of people. As the tourism domain is generally perceived as a medium-low involvement domain (compared to, e.g., shared apartments as an example of high involvement domain [29]), future work should study the perceived importance of privacy risk and disclosure benefit in domains that have higher levels of involvement and/or risk.

User-tailored privacy for group explanations. User-tailored privacy has been proposed and studied as a human-centric solution to reduce users' privacy concerns using recommender systems [57, 62]. As suggested by advocates of User-Tailored Privacy, it

makes it easier to manage one's privacy by automatically tailoring a system's privacy settings to the user's preferences [57]. Future work can utilize the findings of this thesis to automatically predict a balance between users' desire for privacy and their need for transparency. Predicting this balance can facilitate group decision-making, rather than leaving the decision to decide what information they want to disclose as a recurring burden on the users themselves.

Group modeling. In this thesis, we saw a significant difference between the privacy/disclosure preference of different people (e.g., depending on their personality or whether their preferences were in the minority or majority). These individual differences may result in situations where the availability of information is asymmetric (e.g., one user wants to hide their location while the other two users disclose it). Future work should leverage existing work on preference aggregation strategies (e.g., [30, 75]) to address the challenge of reconciling these differences in privacy/disclosure preferences when generating explanations for the entire group. This thesis should ultimately lead to the automatic generation of privacy-preserving explanations for group recommendations adapted to the individual and situational factors identified in our experiments.

6.4 Summary

In this thesis, we designed explanations for group recommendations to increase group members' satisfaction with the recommended items. We started by studying different social choice-based aggregation strategies as a basis to generate group recommendations first to develop a general explanation for the recommended items. We then personalized explanations by adding more personal information aiming to help group members make informed decisions based on each other's preferences/needs. The user comments highlight the need to consider privacy in this context. While we initially considered explaining the aggregation strategies as central, we discovered that other factors such as privacy and group composition (e.g., whether the user's preferences align with the majority in the group, their social relations in the group) were much more influential in the settings that we studied. Upon further analysis, we found that group members' disclosure decision depends on their personality traits (extroversion, agreeableness, and conscientiousness), preference scenario, relationship type, general privacy concern, trust in the group, perception of privacy risk, and disclosure benefit. This resulted in a privacy disclosure model containing different individual and situational models/characteristics that help to predict users' disclosure intention in a group decision/recommendation context. The resulting privacy disclosure model can be utilized to customize and adapt group explanations to the user's disclosure preferences (i.e., based on individuals' trade-off between benefit versus risk to decide on information disclosure). For example, when a group member's preferences are not aligned with the preferences of the majority in the group, they perceive more privacy risk and might be willing to disclose less personal information to the group members. Or when the task of finding a suitable destination in the group is framed cooperatively when asking people for personal information, people are less likely to perceive risk. In this case, they might be willing to disclose more personal information to the group members.

This thesis represents a step towards developing explanations for group recommendation/decision systems by taking group members' privacy concerns into consideration. I have also provided an open-source web-based chat-bot (called Tourybot) to create realistic scenarios of group decision-making where users can control the amount of information disclosed when discussing a recommended item to the group. The Tourybot and our study setups can inspire and assist the research community in conducting human-centered experiments in a group recommendation/decision context. The code is available at the following address: <https://osf.io/z3hnp/>. I believe that this thesis will potentially aid researchers in further exploring many aspects of designing explanations for groups. Besides, utilizing and studying the identified privacy disclosure model to customize group explanations remains for future work. Moreover, this should be used for the automatic generation of privacy-preserving explanations for group recommendations, adapted to all identified individual (individual models) and situational factors identified in our experiments. Last but not least, designing explanations for groups should not follow one-size-fits-all explanation formulation approaches. However, they should be customized with human-centered approaches considering different individual and situational characteristics.

Bibliography

- [1] Anne Anastasi and Susana Urbina. *Psychological testing*. Prentice Hall/Pearson Education, 1997.
- [2] Kenneth J Arrow. A difficulty in the concept of social welfare. *Journal of political economy*, 58(4):328–346, 1950.
- [3] Solomon E Asch. Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological monographs: General and applied*, 70(9):1, 1956.
- [4] Gaurav Bansal, David Gefen, et al. The impact of personal dispositions on information sensitivity, privacy concern and trust in disclosing health information online. *Decision support systems*, 49(2):138–150, 2010.
- [5] Gaurav Bansal, Fatemeh Mariam Zahedi, and David Gefen. Do context and personality matter? trust and privacy concerns in disclosing private information online. *Information & Management*, 53(1):1–21, 2016.
- [6] Francesco Barile, Judith Masthoff, and Silvia Rossi. The adaptation of an individual's satisfaction to group context: the role of ties strength and conflicts. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, pages 357–358, 2017.
- [7] Francesco Barile, Shabnam Najafian, Tim Draws, Oana Inel, Alisa Rieger, Rishav Hada, and Nava Tintarev. Toward benchmarking group explanations: Evaluating the effect of aggregation strategies versus explanation. 2021.
- [8] Joeran Beel, Corinna Breitinger, Stefan Langer, Andreas Lommatzsch, and Bela Gipp. Towards reproducibility in recommender-systems research. *User modeling and user-adapted interaction*, 26(1):69–101, 2016.
- [9] Aylin Caliskan Islam, Jonathan Walsh, and Rachel Greenstadt. 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*, pages 35–46, 2014.
- [10] Da Cao, Xiangnan He, Lianhai Miao, Yahui An, Chao Yang, and Richang Hong. Attentive group recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 645–654, 2018.

- [11] An-Jung Cheng, Yan-Ying Chen, Yen-Ta Huang, Winston H Hsu, and Hong-Yuan Mark Liao. Personalized travel recommendation by mining people attributes from community-contributed photos. In *Proceedings of the 19th ACM international conference on Multimedia*, pages 83–92, 2011.
- [12] Ben Choi, Yi Wu, Jie Yu, and Lesley Pek Wee Land. Love at first sight: The interplay between privacy dispositions and privacy calculus in online social connectivity management. *Journal of the Association for Information Systems*, 19(3):4, 2018.
- [13] Sunny Consolvo, Ian E Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 81–90, 2005.
- [14] Paul T Costa and Robert R McCrae. *Neo personality inventory-revised (NEO PI-R)*. Psychological Assessment Resources Odessa, FL, 1992.
- [15] Paul T Costa Jr and Robert R McCrae. Four ways five factors are basic. *Personality and individual differences*, 13(6):653–665, 1992.
- [16] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction*, 18(5):455, 2008.
- [17] Mary J Culnan. ” how did they get my name?”: an exploratory investigation of consumer attitudes toward secondary information use. *MIS quarterly*, pages 341–363, 1993.
- [18] Mary J Culnan and Pamela K Armstrong. Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization science*, 10(1):104–115, 1999.
- [19] Patrick Dattalo. *Determining sample size: Balancing power, precision, and practicality*. oxford university press, 2008.
- [20] Amra Delic, Judith Masthoff, Julia Neidhardt, and Hannes Werthner. How to use social relationships in group recommenders: empirical evidence. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*, pages 121–129, 2018.
- [21] Amra Delic, Judith Masthoff, and Hannes Werthner. The effects of group diversity in group decision-making process in the travel and tourism domain. In *Information and Communication Technologies in Tourism 2020*, pages 117–129. Springer, 2020.
- [22] Tamara Dinev, Massimo Bellotto, Paul Hart, Vincenzo Russo, Ilaria Serra, and Christian Colautti. Privacy calculus model in e-commerce—a study of italy and the united states. *European Journal of Information Systems*, 15(4):389–402, 2006.

-
- [23] Anne Dohrenwend. Serving up the feedback sandwich. *Family practice management*, 9(10):43, 2002.
- [24] Paul Ekman and Wallace V Friesen. A new pan-cultural facial expression of emotion. *Motivation and emotion*, 10(2):159–168, 1986.
- [25] M Max Evans, Ilja Frissen, and Chun Wei Choo. 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.
- [26] F Faul, E Erdfelder, AG Lang, and A Buchner. A flexible statistical power analysis program for the social, behavioral and biomedical sciences. *Behavior Research Methods*.
- [27] Franz Faul, Edgar Erdfelder, Albert Georg Lang, and Axel Buchner. G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2):175–191, 2007. ISSN 1554351X. doi: 10.3758/BF03193146.
- [28] Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner. G* power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods*, 39(2):175–191, 2007.
- [29] Alexander Felfernig, Muesluem Atas, Thi Ngoc Trang Tran, Martin Stettinger, Seda Polat Erdeniz, and Gerhard Leitner. An analysis of group recommendation heuristics for high-and low-involvement items. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*, pages 335–344. Springer, 2017.
- [30] Alexander Felfernig, Ludovico Boratto, Martin Stettinger, and Marko Tkalčič. Explanations for groups. In *Group Recommender Systems*, pages 105–126. Springer, 2018.
- [31] Shanshan Feng and Jian Cao. Improving group recommendations via detecting comprehensive correlative information. *Multimedia Tools and Applications*, 76(1): 1355–1377, 2017.
- [32] Donelson R Forsyth. *Group dynamics*. Cengage Learning, 2018.
- [33] Mike Gartrell, Xinyu Xing, Qin Lv, Aaron Beach, Richard Han, Shivakant Mishra, and Karim Seada. Enhancing group recommendation by incorporating social relationship interactions. In *Proceedings of the 16th ACM international conference on Supporting group work*, pages 97–106, 2010.
- [34] Albert Gatt and Ehud Reiter. Simplenlg: A realisation engine for practical applications. In *Proceedings of the 12th European Workshop on Natural Language Generation (ENLG 2009)*, pages 90–93, 2009.

- [35] Fatih Gedikli, Dietmar Jannach, and Mouzhi Ge. How should i explain? a comparison of different explanation types for recommender systems. *International Journal of Human-Computer Studies*, 72(4):367–382, 2014.
- [36] Lewis R Goldberg. The development of markers for the big-five factor structure. *Psychological assessment*, 4(1):26, 1992.
- [37] Steven M Graham, Julie Y Huang, Margaret S Clark, and Vicki S Helgeson. The positives of negative emotions: Willingness to express negative emotions promotes relationships. *Personality and Social Psychology Bulletin*, 34(3):394–406, 2008.
- [38] Mark S Granovetter. The strength of weak ties. *American journal of sociology*, 78(6): 1360–1380, 1973.
- [39] Peter Gräsch, Alexander Felfernig, and Florian Reinfrank. Recommend: Towards critiquing-based recommendation with speech interaction. In *Proceedings of the 7th ACM Conference on Recommender Systems*, pages 157–164, 2013.
- [40] Frank E Harrell. Ordinal logistic regression. In *Regression modeling strategies*, pages 311–325. Springer, 2015.
- [41] William J Havlena and Wayne S DeSarbo. On the measurement of perceived consumer risk. *Decision Sciences*, 22(4):927–939, 1991.
- [42] Jonathan L Herlocker, Joseph A Konstan, and John Riedl. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, pages 241–250. ACM, 2000.
- [43] Daniel Herzog and Wolfgang Wörndl. 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*, pages 130–138. ACM, 2019.
- [44] Daniel Herzog and Wolfgang Wörndl. User-centered evaluation of strategies for recommending sequences of points of interest to groups. In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 96–100, 2019.
- [45] David A Hofmann. Overcoming the obstacles to cross-functional decision making: laying the groundwork for collaborative problem solving. *Organizational Dynamics*, 44(1):17–25, 2015.
- [46] Athina Ioannou, Iis Tussyadiah, and Graham Miller. That’s private! understanding travelers’ privacy concerns and online data disclosure. *Journal of Travel Research*, page 0047287520951642, 2020.
- [47] Anthony Jameson, Stephan Baldes, and Thomas Kleinbauer. Enhancing mutual awareness in group recommender systems. In *Proceedings of the IJCAI*, 2003.
- [48] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. A survey on conversational recommender systems. *arXiv preprint arXiv:2004.00646*, 2020.

-
- [49] Oliver P John, Sanjay Srivastava, et al. The big five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research*, 2(1999):102–138, 1999.
- [50] Adam N Joinson, Ulf-Dietrich Reips, Tom Buchanan, and Carina B Paine Schofield. Privacy, trust, and self-disclosure online. *Human-Computer Interaction*, 25(1):1–24, 2010.
- [51] Mohsen Jozani, Emmanuel Ayaburi, Myung Ko, and Kim-Kwang Raymond Choo. Privacy concerns and benefits of engagement with social media-enabled apps: A privacy calculus perspective. *Computers in Human Behavior*, 107:106260, 2020.
- [52] Iris A Junglas, Norman A Johnson, and Christiane Spitzmüller. Personality traits and concern for privacy: an empirical study in the context of location-based services. *European Journal of Information Systems*, 17(4):387–402, 2008.
- [53] Öykü Kapcak, Simone Spagnoli, Vincent Robbmond, Soumitri Vadali, Shabnam Najafian, and Nava Tintarev. Tourexplain: A crowdsourcing pipeline for generating explanations for groups of tourists. In *Workshop on Recommenders in Tourism co-located with the 12th ACM Conference on Recommender Systems (RecSys 2018)*, volume 2222. CEUR, 2018.
- [54] Flavius Kehr, Tobias Kowatsch, Daniel Wentzel, and Elgar Fleisch. Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal*, 25(6):607–635, 2015.
- [55] Mark J Keith, Samuel C Thompson, Joanne Hale, Paul Benjamin Lowry, and Chapman Greer. Information disclosure on mobile devices: Re-examining privacy calculus with actual user behavior. *International journal of human-computer studies*, 71(12):1163–1173, 2013.
- [56] Dongyeon Kim, Kyuhong Park, Yongjin Park, and Jae-Hyeon Ahn. Willingness to provide personal information: Perspective of privacy calculus in iot services. *Computers in Human Behavior*, 92:273–281, 2019.
- [57] Bart P Knijnenburg. Privacy? i can’t even! making a case for user-tailored privacy. *IEEE Security & Privacy*, 15(4):62–67, 2017.
- [58] Bart P Knijnenburg and Martijn C Willemsen. Evaluating recommender systems with user experiments. In *Recommender Systems Handbook*, pages 309–352. Springer, 2015.
- [59] Bart P Knijnenburg, Niels JM Reijmer, and Martijn C Willemsen. Each to his own: how different users call for different interaction methods in recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, pages 141–148, 2011.
- [60] Bart P Knijnenburg, Martijn C Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*, 22(4-5):441–504, 2012.

- [61] Bart P Knijnenburg, Alfred Kobsa, and Hongxia Jin. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies*, 71(12):1144–1162, 2013.
- [62] Bart Piet Knijnenburg. *A user-tailored approach to privacy decision support*. PhD thesis, UC Irvine, 2015.
- [63] Bart Piet Knijnenburg and Alfred Kobsa. Increasing sharing tendency without reducing satisfaction: Finding the best privacy-settings user interface for social networks. In *ICIS*, 2014.
- [64] Joseph A Konstan, Bradley N Miller, David Maltz, Jonathan L Herlocker, Lee R Gordon, and John Riedl. Grouplens: applying collaborative filtering to usenet news. *Communications of the ACM*, 40(3):77–87, 1997.
- [65] Melinda L Korzaan and Katherine T Boswell. The influence of personality traits and information privacy concerns on behavioral intentions. *Journal of Computer Information Systems*, 48(4):15–24, 2008.
- [66] Hanna Krasnova, Sarah Spiekermann, Ksenia Koroleva, and Thomas Hildebrand. Online social networks: Why we disclose. *Journal of information technology*, 25(2): 109–125, 2010.
- [67] S Pavan Kumar and Shilpi Saha. Influence of trust and participation in decision making on employee attitudes in indian public sector undertakings. *Sage Open*, 7(3):2158244017733030, 2017.
- [68] Saurabh Kumar, Pradeep Kumar, and Bharat Bhasker. Interplay between trust, information privacy concerns and behavioural intention of users on online social networks. *Behaviour & Information Technology*, 37(6):622–633, 2018.
- [69] Liselotte Kweekel, Trudie Gerrits, Marlies Rijnders, and Patrick Brown. The role of trust in centeringpregnancy: building interpersonal trust relationships in group-based prenatal care in the netherlands. *Birth*, 44(1):41–47, 2017.
- [70] Robert S Laufer and Maxine Wolfe. Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of social Issues*, 33(3):22–42, 1977.
- [71] Jonathan Lazar. Feng. jh; hochheiser, h. research methods in human-computer interaction, 2010.
- [72] Kyumin Lee, James Caverlee, and Steve Webb. The social honeypot project: protecting online communities from spammers. In *Proceedings of the 19th international conference on World wide web*, pages 1139–1140. ACM, 2010.
- [73] Naresh K Malhotra, Sung S Kim, and James Agarwal. Internet users’ information privacy concerns (iuipc): The construct, the scale, and a causal model. *Information systems research*, 15(4):336–355, 2004.

-
- [74] Judith Masthoff. Group modeling: Selecting a sequence of television items to suit a group of viewers. In *Personalized digital television*, pages 93–141. Springer, 2004.
- [75] Judith Masthoff. Group recommender systems: Combining individual models. In *Recommender systems handbook*, pages 677–702. Springer, 2011.
- [76] Judith Masthoff. Group recommender systems: aggregation, satisfaction and group attributes. In *recommender systems handbook*, pages 743–776. Springer, 2015.
- [77] Judith Masthoff and Albert Gatt. In pursuit of satisfaction and the prevention of embarrassment: affective state in group recommender systems. *User Modeling and User-Adapted Interaction*, 16(3-4):281–319, 2006.
- [78] Roger C Mayer, James H Davis, and F David Schoorman. An integrative model of organizational trust. *Academy of management review*, 20(3):709–734, 1995.
- [79] Kevin McCarthy, Maria Salamó, Lorcan Coyle, Lorraine McGinty, Barry Smyth, and Paddy Nixon. Group recommender systems: a critiquing based approach. In *Proceedings of the 11th international conference on Intelligent user interfaces*, pages 267–269, 2006.
- [80] AKM Mehdy, Michael D Ekstrand, Bart P Knijnenburg, and Hoda Mehrpouyan. Privacy as a planned behavior: Effects of situational factors on privacy perceptions and plans. *arXiv preprint arXiv:2104.11847*, 2021.
- [81] George R Milne and Mary Ellen Gordon. Direct mail privacy-efficiency trade-offs within an implied social contract framework. *Journal of Public Policy & Marketing*, 12(2):206–215, 1993.
- [82] Soha A Mohamed, Taysir Hassan A Soliman, and Adel A Sewisy. A context-aware recommender system for personalized places in mobile applications. *Int. J. Adv. Comput. Sci. Appl*, 7(3):442–448, 2016.
- [83] Shabnam Najafian and Nava Tintarev. Generating consensus explanations for group recommendations. In *UMAP Latebreaking results*, 2018.
- [84] Shabnam Najafian and Nava Tintarev. Generating consensus explanations for group recommendations: an exploratory study. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*, pages 245–250. ACM, 2018.
- [85] Shabnam Najafian, Daniel Herzog, Sihang Qiu, Oana Inel, and Nava Tintarev. You do not decide for me! Evaluating explainable group aggregation strategies for tourism. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, pages 187–196, 2020.
- [86] Shabnam Najafian, Oana Inel, and Nava Tintarev. 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*, pages 85–86, 2020.

- [87] Shabnam Najafian, Amra Delic, Marko Tkalcic, and Nava Tintarev. Factors influencing privacy concern for explanations of group recommendation. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*, pages 14–23, 2021.
- [88] Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. Exploring user concerns about disclosing location and emotion information in group recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, pages 155–164, 2021.
- [89] A. Napierala, M. What Is the Bonferroni correction?, 2012. URL <http://www.aaos.org/news/aaosnow/apr12/research7.asp>.
- [90] Lili Nemeč Zlatolas, Tatjana Welzer, Marko Hölbl, Marjan Heričko, and Aida Kamišalić. A model of perception of privacy, trust, and self-disclosure on online social networks. *Entropy*, 21(8):772, 2019.
- [91] Thuy Ngoc Nguyen. Conversational group recommender systems. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, pages 331–334, 2017.
- [92] Thuy Ngoc Nguyen and Francesco Ricci. Dynamic elicitation of user preferences in a chat-based group recommender system. In *Proceedings of the Symposium on Applied Computing*, pages 1685–1692, 2017.
- [93] Thuy Ngoc Nguyen, Francesco Ricci, Amra Delic, and Derek Bridge. Conflict resolution in group decision making: insights from a simulation study. *User Modeling and User-Adapted Interaction*, 29(5):895–941, 2019.
- [94] Patricia A Norberg, Daniel R Horne, and David A Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs*, 41(1):100–126, 2007.
- [95] Geoff Norman. Likert scales, levels of measurement and the "laws" of statistics. *Advances in Health Sciences Education*, 15(5):625–632, 2010. ISSN 13824996. doi: 10.1007/s10459-010-9222-y.
- [96] Brian Nosek, Johanna Cohoon, Mallory Kidwell, and Jeffrey Robert Spies. Estimating the Reproducibility of Psychological Science. *Science*, 349(6251):aac47160, 2015. doi: 10.1126/science.aac4716.
- [97] Brian A Nosek, George Alter, George C Banks, Denny Borsboom, Sara D Bowman, Steven J Breckler, Stuart Buck, Christopher D Chambers, Gilbert Chin, Garret Christensen, et al. Promoting an open research culture. *Science*, 348(6242):1422–1425, 2015.
- [98] Eirini Ntoutsis, Kostas Stefanidis, Kjetil Nørnvåg, and Hans-Peter Kriegel. Fast group recommendations by applying user clustering. In *International Conference on Conceptual Modeling*, pages 126–140. Springer, 2012.

-
- [99] Ingrid Nunes and Dietmar Jannach. A systematic review and taxonomy of explanations in decision support and recommender systems. *User Modeling and User-Adapted Interaction*, 27(3-5):393–444, 2017.
- [100] Xinru Page, Bart P Knijnenburg, and Alfred Kobsa. Fyi: communication style preferences underlie differences in location-sharing adoption and usage. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 153–162, 2013.
- [101] Jonathan W Palmer, Joseph P Bailey, and Samer Faraj. The role of intermediaries in the development of trust on the www: The use and prominence of trusted third parties and privacy statements. *Journal of Computer-Mediated Communication*, 5(3):JCMC532, 2000.
- [102] Alexandros Paramythis, Stephan Weibelzahl, and Judith Masthoff. Layered evaluation of interactive adaptive systems: framework and formative methods. *User Modeling and User-Adapted Interaction*, 20(5):383–453, 2010.
- [103] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163, 2017.
- [104] Iryna Pentina, Lixuan Zhang, Hatem Bata, and Ying Chen. Exploring privacy paradox in information-sensitive mobile app adoption: A cross-cultural comparison. *Computers in Human Behavior*, 65:409–419, 2016.
- [105] Richard E Petty, John T Cacioppo, and David Schumann. Central and peripheral routes to advertising effectiveness: The moderating role of involvement. *Journal of consumer research*, 10(2):135–146, 1983.
- [106] Robert Plutchik. 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):344–350, 2001.
- [107] Nivedita Prasad. Evaluating explanations for different relationship strengths. Master’s thesis, 2018.
- [108] Lara Quijano-Sanchez, Christian Sauer, Juan A Recio-Garcia, and Belen Diaz-Agudo. Make it personal: a social explanation system applied to group recommendations. *Expert Systems with Applications*, 76:36–48, 2017.
- [109] Ehud Reiter and Robert Dale. Building applied natural language generation systems. *Natural Language Engineering*, 3(1):57–87, 1997.
- [110] Ehud Reiter and Robert Dale. *Building natural language generation systems*. Cambridge university press, 2000.
- [111] Andrew J Rohm and George R Milne. Just what the doctor ordered: The role of information sensitivity and trust in reducing medical information privacy concern. *Journal of Business Research*, 57(9):1000–1011, 2004.

- [112] Silvia Rossi, Francesco Cervone, and Francesco Barile. An altruistic-based utility function for group recommendation. In *Transactions on Computational Collective Intelligence XXVIII*, pages 25–47. Springer, 2018.
- [113] Jason E Sapp, Dario M Torre, Kelsey L Larsen, Eric S Holmboe, and Steven J Durning. Trust in group decisions: a scoping review. *BMC medical education*, 19(1):1–13, 2019.
- [114] Christophe Senot, Dimitre Kostadinov, Makram Bouzid, Jérôme Picault, Armen Aghasaryan, and Cédric Bernier. Analysis of strategies for building group profiles. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 40–51. Springer, 2010.
- [115] Dong-Hee Shin. The effects of trust, security and privacy in social networking: A security-based approach to understand the pattern of adoption. *Interacting with computers*, 22(5):428–438, 2010.
- [116] Dong-Hee Shin. Ubiquitous computing acceptance model: end user concern about security, privacy and risk. *International Journal of Mobile Communications*, 8(2): 169–186, 2010.
- [117] Rashmi Sinha and Kirsten Swearingen. The role of transparency in recommender systems. In *CHI'02 extended abstracts on Human factors in computing systems*, pages 830–831, 2002.
- [118] H Jeff Smith, Tamara Dinev, and Heng Xu. Information privacy research: an interdisciplinary review. *MIS quarterly*, pages 989–1015, 2011.
- [119] Barry Smyth and Paul McClave. Similarity vs. diversity. In *International conference on case-based reasoning*, pages 347–361. Springer, 2001.
- [120] Sanjay Srivastava, Oliver P John, Samuel D Gosling, and Jeff Potter. Development of personality in early and middle adulthood: Set like plaster or persistent change? *Journal of personality and social psychology*, 84(5):1041, 2003.
- [121] Jacqueline Tanghe, Barbara Wisse, and Henk Van Der Flier. The role of group member affect in the relationship between trust and cooperation. *British Journal of Management*, 21(2):359–374, 2010.
- [122] David G Taylor, Donna F Davis, and Ravi Jillapalli. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electronic commerce research*, 9(3):203–223, 2009.
- [123] Nava Tintarev and Judith Masthoff. Effective explanations of recommendations: user-centered design. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 153–156. ACM, 2007.
- [124] Nava Tintarev and Judith Masthoff. A survey of explanations in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*, pages 801–810. IEEE, 2007.

-
- [125] Nava Tintarev and Judith Masthoff. Over-and underestimation in different product domains. In *Workshop on Recommender Systems associated with ECAL*, pages 14–19. Springer Boston, MA, 2008.
- [126] Nava Tintarev and Judith Masthoff. Evaluating recommender explanations: Problems experienced and lessons learned for evaluation of adaptive systems. In *In the workshop on User-Centred Design and Evaluation of Adaptive Systems in association with UMAP'09*. Citeseer, 2009.
- [127] Nava Tintarev and Judith Masthoff. Evaluating the effectiveness of explanations for recommender systems. *User Modeling and User-Adapted Interaction*, 22(4-5):399–439, 2012.
- [128] Nava Tintarev and Judith Masthoff. Explaining recommendations: Design and evaluation. In *Recommender systems handbook*, pages 353–382. Springer, 2015.
- [129] Claudia Toma and Fabrizio Butera. Cooperation versus competition effects on information sharing and use in group decision-making. *Social and Personality Psychology Compass*, 9(9):455–467, 2015.
- [130] Claudia Toma, Ingrid Gilles, and Fabrizio Butera. Strategic use of preference confirmation in group decision making: The role of competition and dissent. *British Journal of Social Psychology*, 52(1):44–63, 2013.
- [131] Thi Ngoc Trang Tran, Müslüm Atas, Alexander Felfernig, Viet Man Le, Ralph Samer, and Martin Stettinger. Towards social choice-based explanations in group recommender systems. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*, pages 13–21. ACM, 2019.
- [132] Janice Y Tsai, Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Location-sharing technologies: Privacy risks and controls. *Isjlp*, 6:119, 2010.
- [133] Thomas P Van Dyke, Vishal Midha, and Hamid Nemati. The effect of consumer privacy empowerment on trust and privacy concerns in e-commerce. *Electronic Markets*, 17(1):68–81, 2007.
- [134] Xinru Wang and Ming Yin. Are explanations helpful? a comparative study of the effects of explanations in ai-assisted decision-making. In *26th International Conference on Intelligent User Interfaces*, pages 318–328, 2021.
- [135] Zhu Wang, Xingshe Zhou, Zhiwen Yu, Haipeng Wang, and Hongbo Ni. Quantitative evaluation of group user experience in smart spaces. *Cybernetics and Systems: An International Journal*, 41(2):105–122, 2010.
- [136] Feng Xu, Katina Michael, and Xi Chen. Factors affecting privacy disclosure on social network sites: an integrated model. *Electronic Commerce Research*, 13(2):151–168, 2013.
- [137] Heng Xu, Tamara Dinev, H Jeff Smith, and Paul Hart. Examining the formation of individual's privacy concerns: Toward an integrative view. 2008.

-
- [138] Luyan Xu, Xuan Zhou, and Ujwal Gadiraju. 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*, pages 91–100, 2020.
- [139] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, pages 22–32, 2005.

List of Figures

3.1	Screenshots of the participants' task	25
3.2	The initial ratings (range [1,6]) of the active user for the recommended POIs by each strategy.	31
3.3	Number of POIs recommended by none, one, two, three or all four strategies.	32
3.4	Pairwise comparison of frequency of occurrence of all 42 predefined POIs between strategies	32
3.5	The distribution of similarity of travel preferences between users in high divergence groups and low divergence groups	33
3.6	Participants' mean fairness, consensus, and satisfaction perception across explanation types	41
3.7	Average satisfaction, by different explanation types and scenarios. Whisker plot depicts 1 SD.	51
4.1	Screenshot of the system that (dynamically) generates group explanations	62
4.2	The conceptual model of perceived risk of information disclosure in a group explanation	81
4.3	The structured equation modeling (SEM) results of factors influencing privacy risk	82
4.4	A screenshot of the chat between an active user and his group	90
4.5	A screenshot of the chat between the chat-bot and an active user	90
4.6	Overview of the experiment procedure for each participant	93
5.1	A screenshot of the chat between an active user and their group	113
5.2	A screenshot of the chat between an active user and the chat-bot	113
5.3	Conceptual model for antecedents of personal information disclosure.	119
5.4	The Structural Equation Model (SEM) for the data of the experiment	122
5.5	The jittered scatterplot of perceived privacy risk and the corresponding information disclosure	123

List of Tables

2.1	An example of group ratings for some restaurants	9
3.1	An example of the <i>Least+</i> strategy	22
3.2	An example of the <i>Fair+</i> strategy	23
3.3	An example of the <i>Average</i> strategy	23
3.4	An example of the <i>Dictatorship (Most Respected Person)</i> strategy	24
3.5	MANOVA: Wilks Test between the levels of divergence and the interaction between the strategies and the levels of divergence	28
3.6	The average and deviations of the user perceived individual satisfaction, group satisfaction, fairness, and acceptance	29
3.7	Ratings of group members for the restaurants	37
3.8	All possible explanation scenarios that participants saw in our study	38
3.9	Results of three two-way ANOVAs for the dependent variables	42
3.10	Explanation categories and examples.	49
3.11	Possible explanation purposes for groups	52
4.1	Frequencies of the chosen privacy option	67
4.2	Selection of the three privacy options is summarized by eight patterns	67
4.3	Participants' responses to the perceived transparency, privacy threat, local privacy, and satisfaction per scenario	72
4.4	Results of location-related information disclosure from an ordinal logistic regression analyses	96
4.5	Results of emotion-related information disclosure from an ordinal logistic regression analyses	97
4.6	Privacy risk percentage of location and emotion information disclosure	98
5.1	Items used to measure participants' privacy concerns and trust and the corresponding CFA outcome.	116
5.2	The results of the SEM analysis of a privacy disclosure model	121
5.3	Participants' level of disclosure, average disclosure benefit, and average privacy risk per item.	123

Summary

My thesis aims to study what makes good explanations for group recommendations. Have you ever been to lunch with other colleagues on a business trip? Do you recall how long it took you to pick a restaurant? In these situations, group recommender systems could help a group decide, e.g., where to go. A group recommender system is a system that recommends items to groups of users collectively, given their preferences. Imagine you start walking to one restaurant only to discover that Person A wants to eat Halal, Person B has an auto-immune protocol diet, and Person C prefers a low-budget place. After visiting a restaurant, your group might also need to pick where to go next (e.g., a war museum, a cannabis store, etc.). Will you speak out if your preferences do not align with the majority in the group? What about when you do not have close relationships with other group members? Not only is it challenging to cater to multiple preferences, it can also be difficult to surface individual preferences in order to make an informed group decision! Explanations, for such recommendations, in this context, act as complementary information (i.e., in this thesis, this information is textual only). They describe why certain items are recommended to help the group make informed decisions on whether to follow or not follow recommendations. However, there are many types of information to include and many ways to formulate an explanation, and it is unclear which information should be shown in the explanation for a group.

In Chapter 3, we started evaluating with people what information, in general, should be conveyed when generating a group explanation. Similar to explanations for single-user recommendations, explanations for groups can be designed based on the underlying recommendation algorithm, which is aggregation strategies in this thesis (a form of recommendation algorithm for groups that aggregate individual item-ratings predictions which are given). Based on our results, it seems that explanations containing information about the aggregation strategy do not significantly benefit group members (i.e., increase group members' satisfaction). Besides, user comments highlight the need to protect certain types of information, i.e., group members' ratings of items. This suggests that studying other factors, such as privacy (i.e., protecting certain kinds of information from the group), appears to be a more promising direction than explaining aggregation strategies.

Grounded on the findings of Chapter 3, we decided to go deeper into what people would disclose in a group explanation. In Chapter 4, we found that there may be some individual differences in the level of privacy risk. This led us to another experiment where we investigated which factors contribute to these individual differences. We studied some factors identified in the literature that influence individual privacy risk. Our results confirmed that these factors seem to matter in our experiment's context as well, namely: **a**) group members' personality (using the 'Big Five' personality traits), **b**) specific preference scenarios (i.e., whether the user's preferences are aligned or not aligned with the preferences of the majority in the group), **c**) the type of relationship they have in the group (both relationship strength and equality of positions, i.e., loosely coupled heterogeneous, versus

tightly coupled homogeneous). In a follow-up experiment where we studied the effects of these factors on people's actual disclosure decisions, surprisingly, we did not find an effect of personality or preference scenario on group members' disclosure decisions. This result can be explained by the task design, as our study asked users to convince other group members to accept or skip the suggested place. Therefore, one explanation for why we did not obtain the expected results is that our task design nudged participants into a "convincing mindset". Therefore, in Chapter 5, we studied the effect of task design and other potential intermediate factors on disclosure behavior.

In Chapter 5, we found that task design (whether group members were instructed to convince other group members of their opinion or not) affected participants' disclosure decisions (for certain types of sensitive information). The influence of task on actual disclosure leads to the idea that disclosure benefit and privacy risk of information disclosure could cancel each other out in actual disclosure decision and result in smaller effects of the factors mentioned earlier (i.e., personality, etc.) on participants' information disclosure. This leads us to another study where we investigate how people make the trade-off between *disclosure benefit* (i.e., the extent to which users believe disclosing their personal information to their group members is beneficial for the group decision or their own negotiation position within the group) versus *privacy risk* (i.e., an expectation of losses associated with the disclosure of personal information in the group) to deciding on information disclosure in a group explanation. Results indicate that individuals' personality (using the 'Big Five' personality traits) and their preference scenario (whether one's preferences align with the majority in the group) affect their perception of general privacy. This, in turn, affects their trust in the group, which affects their perception of privacy risk and disclosure benefit when disclosing personal information, ultimately influencing the amount of personal information they disclose. Besides, an interesting finding indicates that privacy risk on information disclosure is different for different types of tasks, specifically for a competitive task (i.e., instructed to convince the group to visit or skip a recommended place) versus a cooperative task (i.e., instructed to reach a decision in the group).⁴ Privacy risk significantly impacts information disclosure when the task of finding a suitable destination is framed competitively but not when it is framed cooperatively. This suggests the cooperative formulation of tasks can be used in group recommendations to facilitate reaching consensus.

This thesis findings contribute to a better understanding of the moderating factors of information disclosure in group decision-making and shed new light on the role of task design on information disclosure. This thorough investigation of the dynamics between these factors and disclosure resulted in a user modeling theory (representation of users for the purpose of understanding user's disclosure behavior) that may inform considerations for automatically generating group explanations from a human-centered perspective. We concluded the thesis with design recommendations for developing explanations in group recommendation/decision-making systems.

⁴Note that when designing the experiment, we thought of task design in terms of convincing people as in the task design experiment. However, in the new study design, we saw in user comments that they perceived that task more as cooperative and competitive settings (see Section 5.2).

Samenvatting

Mijn proefschrift heeft tot doel om te onderzoeken wat goede verklaringen zijn voor groepsaanbevelingen. Ben je wel eens met andere collega's gaan lunchen op een zakenreis? Weet je nog hoe lang het duurde voordat je een restaurant uitkoos? In deze situaties kunnen groepsaanbevelingssystemen een groep helpen beslissen, bijvoorbeeld waar ze heen gaan. Een groepsaanbevelingssysteem is een systeem dat items gezamenlijk aanbeveelt aan groepen gebruikers, gegeven hun voorkeuren. Stel je voor dat je naar een restaurant loopt en ontdekt dat persoon A halal wil eten, persoon B een dieet met auto-immuunprotocol heeft en persoon C de voorkeur geeft aan een goedkope plek. Nadat je een restaurant hebt bezocht, moet je groep misschien ook kiezen waar ze heen willen (bijvoorbeeld een oorlogsmuseum, een cannabiswinkel, enz.). Spreek je je uit als je voorkeuren niet overeenkomen met de meerderheid in de groep? Hoe zit het als je geen nauwe relaties hebt met andere groepsleden? Het is niet alleen een uitdaging om tegemoet te komen aan meerdere voorkeuren, het kan ook moeilijk zijn om individuele voorkeuren naar voren te brengen om een weloverwogen groepsbeslissing te nemen! Een verklaring of uitleg voor dergelijke aanbevelingen fungeert in deze context als aanvullende informatie (n.b. in dit proefschrift is deze informatie alleen tekstueel). De uitleg beschrijft waarom bepaalde items worden aanbevolen om de groep te helpen weloverwogen beslissingen te nemen over het al dan niet opvolgen van aanbevelingen. Er zijn echter veel soorten informatie om op te nemen en veel manieren om een verklaring of uitleg te formuleren, en het is onduidelijk welke informatie in de uitleg voor een groep moet worden getoond.

In hoofdstuk 3 zijn we begonnen te evalueren met mensen welke informatie in het algemeen moet worden overgebracht bij het genereren van een groepsuitleg. Net als bij verklaringen voor aanbevelingen voor één gebruiker, kunnen verklaringen voor groepen worden ontworpen op basis van het onderliggende aanbevelingsalgoritme, en in dit proefschrift gebruiken we daarvoor aggregatiestrategieën (een vorm van aanbevelingsalgoritme voor groepen die de individuele voorspellingen van itemscores die worden gegeven aggregeren). Op basis van onze resultaten lijkt het erop dat verklaringen met informatie over de aggregatiestrategie de groepsleden niet significant ten goede komen (d.w.z. de tevredenheid van de groepsleden verhogen). Bovendien benadrukken gebruikerscommentaren de noodzaak om bepaalde soorten informatie te beschermen, d.w.z. de door groepsleden gegeven beoordelingen van items. Dit suggereert dat het bestuderen van andere factoren, zoals privacy (d.w.z. het beschermen van bepaalde soorten informatie van de groep), een meer veelbelovende richting lijkt te zijn dan het uitleggen van aggregatiestrategieën.

Op basis van de bevindingen van hoofdstuk 3 besloten we dieper in te gaan op wat mensen zouden onthullen in een groepsuitleg. In hoofdstuk 4 ontdekten we dat er enkele individuele verschillen kunnen zijn in het niveau van privacyrisico. Dit leidde ons naar een ander experiment waarin we onderzochten welke factoren bijdragen aan deze individuele verschillen. We hebben enkele in de literatuur geïdentificeerde factoren bestu-

deerd die van invloed zijn op het individuele privacyrisico. Onze resultaten bevestigden dat deze factoren ook van belang lijken te zijn in de context van ons experiment, namelijk: a) de persoonlijkheid van de groepsleden (met behulp van de 'Big Five' persoonlijkheidskenmerken), b) specifieke voorkeursscenario's (d.w.z. of de voorkeuren van de gebruiker wel of niet zijn afgestemd op de voorkeuren van de meerderheid in de groep), c) het type relatie dat ze in de groep hebben (zowel relatiesterkte als gelijkheid van posities, d.w.z. losjes gekoppeld heterogeen versus nauw gekoppeld homogeen). In een vervolgonderzoek waarin we de effecten van deze factoren op de feitelijke openbaarmakingsbeslissingen van mensen bestudeerden, vonden we verrassend genoeg geen effect van persoonlijkheid of voorkeursscenario op de openbaarmakingsbeslissingen van groepsleden. Dit resultaat kan worden verklaard door het taakontwerp, aangezien onze studie gebruikers vroeg om andere groepsleden te overtuigen om de voorgestelde plaats te accepteren of over te slaan. Daarom is een van de redenen waarom we niet de geanticiperde resultaten behaalden, dat ons taakontwerp de deelnemers in een 'mindset van overtuiging' duwde. Daarom hebben we in hoofdstuk 5 het effect van taakontwerp en andere mogelijke intermediaire factoren op openbaarmakingsgedrag bestudeerd.

In hoofdstuk 5 ontdekten we dat taakontwerp (of groepsleden de opdracht kregen om andere groepsleden van hun mening te overtuigen of niet) van invloed was op de openbaarmakingsbeslissingen van deelnemers (voor bepaalde soorten gevoelige informatie). De invloed van de taak op feitelijke openbaarmaking leidt tot het idee dat het openbaarmakingsvoordeel en het privacyrisico van openbaarmaking van informatie elkaar zouden kunnen opheffen bij een daadwerkelijke openbaarmakingsbeslissing en zo resulteren in kleinere effecten van de eerder genoemde factoren (d.w.z. persoonlijkheid, enz.) op het vrijgeven van informatie. Dit leidt ons naar een ander onderzoek waarin we onderzoeken hoe mensen de afweging maken tussen openbaarmakingsvoordeel (d.w.z. de mate waarin gebruikers geloven dat het vrijgeven van hun persoonlijke informatie aan hun groepsleden gunstig is voor de groepsbeslissing of hun eigen onderhandelingspositie binnen de groep) versus privacyrisico (d.w.z. een verwachting van verliezen in verband met de openbaarmaking van persoonlijke informatie in de groep) in relatie tot het beslissen over openbaarmaking van informatie in een groepsverklaring. De resultaten geven aan dat de persoonlijkheid van individuen (met behulp van de 'Big Five'-persoonlijkheidskenmerken) en hun voorkeursscenario (of iemands voorkeuren overeenkomen met de meerderheid in de groep) van invloed zijn op hun perceptie van algemene privacy. Dit heeft op zijn beurt invloed op hun vertrouwen in de groep, wat van invloed is op hun perceptie van privacyrisico's en openbaarmakingsvoordeel bij het vrijgeven van persoonlijke informatie, en is zo uiteindelijk van invloed op de hoeveelheid persoonlijke informatie die ze vrijgeven. Bovendien geeft een interessante bevinding aan dat het privacyrisico bij het vrijgeven van informatie verschillend is voor verschillende soorten taken, met name voor een competitieve taak (d.w.z. geïnstrueerd om de groep te overtuigen een aanbevolen plaats te bezoeken of over te slaan) versus een coöperatieve taak (d.w.z. geïnstrueerd om tot een beslissing komen in de groep).⁵ Privacyrisico heeft een significante invloed op het vrijgeven van informatie wanneer de taak om een geschikte bestemming te vinden competitief wordt ingekaderd,

⁵Merk op dat we bij het ontwerpen van het experiment dachten in termen van het overtuigen van mensen, zoals in het taakontwerpexperiment. In de nieuwe onderzoeksopzet zagen we echter in opmerkingen van gebruikers dat ze die taak meer als een coöperatieve en competitieve omgeving zagen (zie paragraaf 5.2).

maar niet wanneer deze in samenwerking wordt vormgegeven. Dit suggereert dat de gezamenlijke formulering van taken kan worden gebruikt in groepsaanbevelingen om het bereiken van consensus te vergemakkelijken.

De bevindingen van dit proefschrift dragen bij aan een beter begrip van de modere-rende factoren van het vrijgeven van informatie in groepsbesluitvorming en werpen een nieuw licht op de rol van taakontwerp bij het vrijgeven van informatie. Dit grondige onderzoek naar de dynamiek tussen deze factoren en het vrijgeven van informatie resul-teerde in een gebruikersmodelleringstheorie (representatie van gebruikers met als doel het gedrag van vrijgeven door gebruikers te begrijpen) die overwegingen kan opleveren voor het automatisch genereren van groepsverklaringen vanuit een mensgericht perspec-tief. We sloten het proefschrift af met ontwerpaanbevelingen voor het ontwikkelen van verklaringen in groepsaanbevelings-/besluitvormingssystemen.

Curriculum Vitæ

Shabnam Najafian was born in Zanzan, Iran. Graduated from the Technical University of Munich (TUM), Germany, with a master's degree. Her dissertation focused on developing and evaluating context-aware user interactions for mobile recommender systems. She worked part-time at Max Planck Digital Library as a data scientist during her master's degree, where she designed and developed RESTful web services for 3D visualization, analysis, and reporting of neuroscientific data in large neuron databases. After, she started working as a software developer at a top software development company (Freiheit.com) in Hamburg, Germany, where she worked on implementing agile project management software by building microservices in Go-lang. Earlier in her career, she graduated from the Zanzan Institute of Advanced Studies in Basic Sciences in Iran with a bachelor's degree. Her first job after graduation was as a database programmer at Asan Pardakht, Tehran, Iran, where she dealt with sensitive bank account data.

From February 2018, Shabnam was a Ph.D. student in the Web Information Systems group at the Delft University of Technology, supervised by prof. dr. Nava Tintarev. In her Ph.D., she designed natural language explanations for group recommendations, taking into account users' privacy concerns by analyzing individual and situational characteristics of group members. To create realistic scenarios of group decision-making, she developed an open-source chatbot for a group chat to discuss recommendations. Shabnam's research has been published in leading conferences and journals in related domains (e.g., ACM RecSys, ACM UMAP, ACM Hypertext, IUI, and the UMUI Journal). As a reviewer for several conferences, she served on the reviews for UMUI, ACM RecSys, and ACM UMAP. She has a particular interest in human-centered computing, behavioral computer science, data analysis, and natural language generation.

Publications

1. Shabnam Najafian and Nava Tintarev. Generating consensus explanations for group recommendations. In *UMAP Latebreaking results*, 2018
2. Öykü Kapcak, Simone Spagnoli, Vincent Robbmond, Soumitri Vadali, Shabnam Najafian, and Nava Tintarev. Tourexplain: A crowdsourcing pipeline for generating explanations for groups of tourists. In *Workshop on Recommenders in Tourism colocated with the 12th ACM Conference on Recommender Systems (RecSys 2018)*, volume 2222. CEUR, 2018
3. Shabnam Najafian, Daniel Herzog, Sihang Qiu, Oana Inel, and Nava Tintarev. You do not decide for me! Evaluating explainable group aggregation strategies for tourism. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, pages 187–196, 2020

4. Shabnam Najafian, Oana Inel, and Nava Tintarev. 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*, pages 85–86, 2020
5. Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. Exploring user concerns about disclosing location and emotion information in group recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, pages 155–164, 2021
6. Francesco Barile, Shabnam Najafian, Tim Draws, Oana Inel, Alisa Rieger, Rishav Hada, and Nava Tintarev. Toward benchmarking group explanations: Evaluating the effect of aggregation strategies versus explanation. 2021
7. Shabnam Najafian, Amra Delic, Marko Tkalcic, and Nava Tintarev. Factors influencing privacy concern for explanations of group recommendation. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*, pages 14–23, 2021
8. Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. Exploring user concerns about disclosing location and emotion information in group recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, pages 155–164, 2021
9. Shabnam Najafian, Geoff Musick, Bart Knijnenburg, and Nava Tintarev. How do People Make Decisions in Disclosing Personal Information in Tourism Group Recommendations in Competitive versus Cooperative Conditions? *Modeling and User-Adapted Interaction (UMUAI) Journal*, 2022 (under review)

SIKS Dissertation Series

Since 1998, all dissertations written by Ph.D. students who have conducted their research under auspices of a senior research fellow of the SIKS research school are published in the SIKS Dissertation Series.

—1998—

- 1998-01** Johan van den Akker, *DEGAS - An Active, Temporal Database of Autonomous Objects*, CWI
1998-02 Floris Wiesman, *Information Retrieval by Graphically Browsing Meta-Information*, UM
1998-03 Ans Steuten, *A Contribution to the Linguistic Analysis of Business Conversations within the Language/Action Perspective*, TUD
1998-04 Dennis Breuker, *Memory versus Search in Games*, UM
1998-05 E.W.Oskamp, *Computerondersteuning bij Straftoemeting*, RUL

—1999—

- 1999-01** Mark Sloof, *Physiology of Quality Change Modelling; Automated modelling of Quality Change of Agricultural Products*, VU
1999-02 Rob Potharst, *Classification Using Decision Trees and Neural Nets*, EUR
1999-03 Don Beal, *The Nature of Minimax Search*, UM
1999-04 Jacques Penders, *The practical Art of Moving Physical Objects*, UM
1999-05 Aldo de Moor, *Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems*, KUB
1999-06 Niek J.E. Wijngaards, *Re-design of Compositional Systems*, VU
1999-07 David Spelt, *Verification Support for Object Database Design*, UT
1999-08 Jacques H.J. Lenting, *Informed Gambling: Conception and Analysis of a Multi-Agent Mechanism for Discrete Reallocation*, UM

—2000—

- 2000-01** Frank Niessink, *Perspectives on Improving Software Maintenance*, VU
2000-02 Koen Holtman, *Prototyping of CMS Storage Management*, TUE
2000-03 Carolien M.T. Metselaar, *Sociaal-organisatorische Gevolgen van Kennistechnologie; een Procesbenadering en Actorperspectief*, UvA

- 2000-04** Geert de Haan, *ETAG, A Formal Model of Competence Knowledge for User Interface Design*, VU
2000-05 Ruud van der Pol, *Knowledge-Based Query Formulation in Information Retrieval*, UM
2000-06 Rogier van Eijk, *Programming Languages for Agent Communication*, UU
2000-07 Niels Peek, *Decision-Theoretic Planning of Clinical Patient Management*, UU
2000-08 Veerle Coup, *Sensitivity Analysis of Decision-Theoretic Networks*, EUR
2000-09 Florian Waas, *Principles of Probabilistic Query Optimization*, CWI
2000-10 Niels Nes, *Image Database Management System Design Considerations, Algorithms, and Architecture*, CWI
2000-11 Jonas Karlsson, *Scalable Distributed Data Structures for Database Management*, CWI

—2001—

- 2001-01** Silja Renooij, *Qualitative Approaches to Quantifying Probabilistic Networks*, UU
2001-02 Koen Hindriks, *Agent Programming Languages: Programming with Mental Models*, UU
2001-03 Maarten van Someren, *Learning as Problem Solving*, UvA
2001-04 Evgueni Smirnov, *Conjunctive and Disjunctive Version Spaces with Instance-Based Boundary Sets*, UM
2001-05 Jacco van Ossenbruggen, *Processing Structured Hypermedia: A Matter of Style*, VU
2001-06 Martijn van Welie, *Task-based User Interface Design*, VU
2001-07 Bastiaan Schonhage, *Diva: Architectural Perspectives on Information Visualization*, VU
2001-08 Pascal van Eck, *A Compositional Semantic Structure for Multi-Agent Systems Dynamics*, VU
2001-09 Pieter Jan 't Hoen, *Towards Distributed Development of Large Object-Oriented Models, Views of Packages as Classes*, RUL
2001-10 Maarten Sierhuis, *Modeling and Simulating Work Practice BRAHMS: a Multi-Agent Modeling and Simulation Language for Work Practice Analysis and Design*, UvA
2001-11 Tom M. van Engers, *Knowledge Management: The Role of Mental Models in Business Systems Design*, VU

—2002—

2002-01 Nico Lassing, *Architecture-Level Modifiability Analysis*, VU

2002-02 Roelof van Zwol, *Modelling and Searching Web-Based Document Collections*, UT

2002-03 Henk Ernst Blok, *Database Optimization Aspects for Information Retrieval*, UT

2002-04 Juan Roberto Castelo Valdueza, *The Discrete Acyclic Digraph Markov Model in Data Mining*, UU

2002-05 Radu Serban, *The Private Cyberspace Modeling Electronic Environments inhabited by Privacy-concerned Agents*, VU

2002-06 Laurens Mommers, *Applied Legal Epistemology; Building a Knowledge-Based Ontology of the Legal Domain*, UL

2002-07 Peter Boncz, *Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications*, CWI

2002-08 Jaap Gordijn *Value Based Requirements Engineering: Exploring Innovative E-Commerce Ideas*, CWI

2002-09 Willem-Jan van den Heuvel, *Integrating Modern Business Applications with Objectified Legacy Systems*, KUB

2002-10 Brian Sheppard *Towards Perfect Play of Scrabble*, UM

2002-11 Wouter C.A. Wijngaards, *Agent Based Modelling of Dynamics: Biological and Organisational Applications*, VU

2002-12 Albrecht Schmidt, *Processing XML in Database Systems*, UvA

2002-13 Hongjing Wu, *A Reference Architecture for Adaptive Hypermedia Applications*, TUE

2002-14 Wieke de Vries, *Agent Interaction: Abstract Approaches to Modelling, Programming and Verifying Multi-Agent Systems*, UU

2002-15 Rik Eshuis, *Semantics and Verification of UML Activity Diagrams for Workflow Modelling*, UT

2002-16 Pieter van Langen, *The Anatomy of Design: Foundations, Models and Applications*, VU

2002-17 Stefan Manegold, *Understanding, Modeling, and Improving Main-Memory Database Performance*, UvA

—2003—

2003-01 Heiner Stuckenschmidt, *Ontology-Based Information Sharing in Weakly Structured Environments*, VU

2003-02 Jan Broersen, *Modal Action Logics for Reasoning About Reactive Systems*, VU

2003-03 Martijn Schuemie, *Human-Computer Interaction and Presence in Virtual Reality Exposure Therapy*, TUD

2003-04 Milan Petkovic, *Content-Based Video Retrieval Supported by Database Technology*, UT

2003-05 Jos Lehmann, *Causation in Artificial Intelligence and Law - A modelling approach*, UvA

2003-06 Boris van Schooten, *Development and Specification of Virtual Environments*, UT

2003-07 Machiel Jansen, *Formal Explorations of Knowledge Intensive Tasks*, UvA

2003-08 Yongping Ran, *Repair Based Scheduling*, UM

2003-09 Rens Kortmann, *The Resolution of Visually Guided Behaviour*, UM

2003-10 Andreas Lincke, *Electronic Business Negotiation: Some Experimental Studies on the Interaction Between Medium, Innovation Context and Culture*, UvT

2003-11 Simon Keizer, *Reasoning Under Uncertainty in Natural Language Dialogue Using Bayesian Networks*, UT

2003-12 Roelend Ordelman, *Dutch Speech Recognition in Multi-Media Information Retrieval*, UT

2003-13 Jeroen Donkers, *Nosce Hostem - Searching with Opponent Models*, UM

2003-14 Stijn Hoppenbrouwers, *Freezing Language: Conceptualisation Processes across ICT-Supported Organisations*, KUN

2003-15 Mathijs de Weerd, *Plan Merging in Multi-Agent Systems*, TUD

2003-16 Menzo Windhouwer, *Feature Grammar Systems - Incremental Maintenance of Indexes to Digital Media Warehouses*, CWI

2003-17 David Jansen, *Extensions of Statecharts with Probability, Time, and Stochastic Timing*, UT

2003-18 Levente Kocsis, *Learning Search Decisions*, UM

—2004—

2004-01 Virginia Dignum, *A Model for Organizational Interaction: Based on Agents, Founded in Logic*, UU

2004-02 Lai Xu, *Monitoring Multi-party Contracts for E-business*, UvT

2004-03 Perry Groot, *A Theoretical and Empirical Analysis of Approximation in Symbolic Problem Solving*, VU

2004-04 Chris van Aart, *Organizational Principles for Multi-Agent Architectures*, UvA

2004-05 Viara Popova, *Knowledge Discovery and Monotonicity*, EUR

2004-06 Bart-Jan Hommes, *The Evaluation of Business Process Modeling Techniques*, TUD

2004-07 Elise Boltjes, *Voorbeeldig Onderwijs; Voorbeeldgestuurd Onderwijs, een Opstap naar Abstract Denken, Vooral voor Meisjes*, UM

2004-08 Joop Verbeek, *Politie en de Nieuwe Internationale Informatiemarkt: Grensregionale Politie Gegevensuitwisseling en Digitale Expertise*, UM

2004-09 Martin Caminada, *For the Sake of the Argument; Explorations into Argument-Based Reasoning*, VU

2004-10 Suzanne Kabel, *Knowledge-rich Indexing of Learning-Objects*, UvA

2004-11 Michel Klein, *Change Management for Distributed Ontologies*, VU

2004-12 The Duy Bui, *Creating Emotions and Facial Expressions for Embodied Agents*, UT

2004-13 Wojciech Jamroga, *Using Multiple Models of Reality: On Agents who Know how to Play*, UT

2004-14 Paul Harrenstein, *Logic in Conflict. Logical Explorations in Strategic Equilibrium*, UU

2004-15 Arno Knobbe, *Multi-Relational Data Mining*, UU

2004-16 Federico Divina, *Hybrid Genetic Relational Search for Inductive Learning*, UvA

2004-17 Mark Winands, *Informed Search in Complex Games*, UM

2004-18 Vania Bessa Machado, *Supporting the Construction of Qualitative Knowledge Models*, UvA

2004-19 Thijs Westerveld, *Using Generative Probabilistic Models for Multi-Media Retrieval*, UT

2004-20 Madelon Evers, *Learning from Design: Facilitating Multidisciplinary Design Teams*, Nyenrode

—2005—

2005-01 Floor Verdenius, *Methodological Aspects of Designing Induction-Based Applications*, UvA

2005-02 Erik van der Werf, *AI Techniques for the Game of Go*, UM

2005-03 Franc Grootjen, *A Pragmatic Approach to the Conceptualisation of Language*, RUN

2005-04 Nirvana Meratnia, *Towards Database Support for Moving Object data*, UT

2005-05 Gabriel Infante-Lopez, *Two-Level Probabilistic Grammar for Natural Language Parsing*, UvA

2005-06 Pieter Spronck, *Adaptive Game AI*, UM

2005-07 Flavius Frasinca, *Hypermedia Presentation Generation for Semantic Web Information Systems*, TUE

2005-08 Richard Vdovjak, *A Model-Driven Approach for Building Distributed Ontology-Based Web Applications*, TUE

2005-09 Jeen Broekstra, *Storage, Querying and Inferencing for Semantic Web Languages*, VU

2005-10 Anders Bouwer, *Explaining Behaviour: Using Qualitative Simulation in Interactive Learning Environments*, UvA

2005-11 Elth Ogston, *Agent Based Matchmaking and Clustering - A Decentralized Approach to Search*, VU

2005-12 Csaba Boer, *Distributed Simulation in Industry*, EUR

2005-13 Fred Hamburg, *Een Computermodel voor het Onderstellen van Euthanasiebeslissingen*, UL

2005-14 Borys Omelayenko, *Web-Service Configuration on the Semantic Web; Exploring how Semantics Meets Pragmatics*, VU

2005-15 Tibor Bosse, *Analysis of the Dynamics of Cognitive Processes*, VU

2005-16 Joris Graaumanns, *Usability of XML Query Languages*, UU

2005-17 Boris Shishkov, *Software Specification Based on Reusable Business Components*, TUD

2005-18 Danielle Sent, *Test-Selection Strategies for Probabilistic Networks*, UU

2005-19 Michel van Dartel, *Situated Representation*, UM

2005-20 Cristina Coteanu, *Cyber Consumer Law, State of the Art and Perspectives*, UL

2005-21 Wijnand Derks, *Improving Concurrency and Recovery in Database Systems by Exploiting Application Semantics*, UT

—2006—

2006-01 Samuil Angelov, *Foundations of B2B Electronic Contracting*, TUE

2006-02 Cristina Chisalita, *Contextual Issues in the Design and Use of Information Technology in Organizations*, VU

2006-03 Noor Christoph, *The Role of Metacognitive Skills in Learning to Solve Problems*, UvA

2006-04 Marta Sabou, *Building Web Service Ontologies*, VU

2006-05 Cees Pierik, *Validation Techniques for Object-Oriented Proof Outlines*, UU

2006-06 Ziv Baida, *Software-Aided Service Bundling - Intelligent Methods & Tools for Graphical Service Modeling*, VU

2006-07 Marko Smiljanic, *XML Schema Matching - Balancing Efficiency and Effectiveness by Means of Clustering*, UT

2006-08 Eelco Herder, *Forward, Back and Home Again - Analyzing User Behavior on the Web*, UT

2006-09 Mohamed Wahdan, *Automatic Formulation of the Auditor's Opinion*, UM

2006-10 Ronny Siebes, *Semantic Routing in Peer-to-Peer Systems*, VU

2006-11 Joeri van Ruth, *Flattening Queries over Nested Data Types*, UT

2006-12 Bert Bongers, *Interactivation - Towards an E-cology of People, our Technological Environment, and the Arts*, VU

2006-13 Henk-Jan Lebbink, *Dialogue and Decision Games for Information Exchanging Agents*, UU

2006-14 Johan Hoorn, *Software Requirements: Update, Upgrade, Re-design - Towards a Theory of Requirements Change*, VU

2006-15 Rainer Malik, *CONAN: Text Mining in the Biomedical Domain*, UU

2006-16 Carsten Riggelsen, *Approximation Methods for Efficient Learning of Bayesian Networks*, UU

2006-17 Stacey Nagata, *User Assistance for Multitasking with Interruptions on a Mobile Device*, UU

2006-18 Valentin Zhizhkhun, *Graph Transformation for Natural Language Processing*, UvA

2006-19 Birna van Riemsdijk, *Cognitive Agent Programming: A Semantic Approach*, UU

2006-20 Marina Velikova, *Monotone Models for Prediction in Data Mining*, UvT

2006-21 Bas van Gils, *Aptness on the Web*, RUN

2006-22 Paul de Vrieze, *Fundamentals of Adaptive Personalisation*, RUN

2006-23 Ion Juvina, *Development of Cognitive Model for Navigating on the Web*, UU

2006-24 Laura Hollink, *Semantic Annotation for Retrieval of Visual Resources*, VU

2006-25 Madalina Drugan, *Conditional Log-Likelihood MDL and Evolutionary MCMC*, UU

2006-26 Vojkan Mihajlovic, *Score Region Algebra: A Flexible Framework for Structured Information Retrieval*, UT

2006-27 Stefano Bocconi, *Vox Populi: Generating Video Documentaries from Semantically Annotated Media Repositories*, CWI

2006-28 Borkur Sigurbjornsson, *Focused Information Access using XML Element Retrieval*, UvA

—2007—

2007-01 Kees Leune, *Access Control and Service-Oriented Architectures*, UvT

2007-02 Wouter Teepe, *Reconciling Information Exchange and Confidentiality: A Formal Approach*, RUG

2007-03 Peter Mika, *Social Networks and the Semantic Web*, VU

2007-04 Jurriaan van Diggelen, *Achieving Semantic Interoperability in Multi-Agent Systems: a Dialogue-Based Approach*, UU

2007-05 Bart Schermer, *Software Agents, Surveillance, and the Right to Privacy: a Legislative Framework for Agent-Enabled Surveillance*, UL

2007-06 Gilad Mishne, *Applied Text Analytics for Blogs*, UvA

- 2007-07 Natasa Jovanovic, *To Whom It May Concern - Addressee Identification in Face-to-Face Meetings*, UT
- 2007-08 Mark Hoogendoorn, *Modeling of Change in Multi-Agent Organizations*, VU
- 2007-09 David Mobach, *Agent-Based Mediated Service Negotiation*, VU
- 2007-10 Huib Aldewereld, *Autonomy vs. Conformity: an Institutional Perspective on Norms and Protocols*, UU
- 2007-11 Natalia Stash, *Incorporating Cognitive/Learning Styles in a General-Purpose Adaptive Hypermedia System*, TUE
- 2007-12 Marcel van Gerven, *Bayesian Networks for Clinical Decision Support: A Rational Approach to Dynamic Decision-Making under Uncertainty*, RUN
- 2007-13 Rutger Rienks, *Meetings in Smart Environments; Implications of Progressing Technology*, UT
- 2007-14 Niek Bergboer, *Context-Based Image Analysis*, UM
- 2007-15 Joyca Lacroix, *NIM: a Situated Computational Memory Model*, UM
- 2007-16 Davide Grossi, *Designing Invisible Handcuffs. Formal Investigations in Institutions and Organizations for Multi-Agent Systems*, UU
- 2007-17 Theodore Charitos, *Reasoning with Dynamic Networks in Practice*, UU
- 2007-18 Bart Orriens, *On the Development and Management of Adaptive Business Collaborations*, UvT
- 2007-19 David Levy, *Intimate Relationships with Artificial Partners*, UM
- 2007-20 Slinger Jansen, *Customer Configuration Updating in a Software Supply Network*, UU
- 2007-21 Karianne Vermaas, *Fast Diffusion and Broadening Use: A Research on Residential Adoption and Usage of Broadband Internet in the Netherlands Between 2001 and 2005*, UU
- 2007-22 Zlatko Zlatev, *Goal-Oriented Design of Value and Process Models from Patterns*, UT
- 2007-23 Peter Barna, *Specification of Application Logic in Web Information Systems*, TUE
- 2007-24 Georgina Ramirez Camps, *Structural Features in XML Retrieval*, CWI
- 2007-25 Joost Schalken, *Empirical Investigations in Software Process Improvement*, VU
- 2008—
- 2008-01 Katalin Boer-Sorbon, *Agent-Based Simulation of Financial Markets: A Modular, Continuous-Time Approach*, EUR
- 2008-02 Alexei Sharpanskykh, *On Computer-Aided Methods for Modeling and Analysis of Organizations*, VU
- 2008-03 Vera Hollink, *Optimizing Hierarchical Menus: a Usage-Based Approach*, UvA
- 2008-04 Ander de Keijzer, *Management of Uncertain Data - Towards Unattended Integration*, UT
- 2008-05 Bela Mutschler, *Modeling and Simulating Causal Dependencies on Process-Aware Information Systems from a Cost Perspective*, UT
- 2008-06 Arjen Hommersom, *On the Application of Formal Methods to Clinical Guidelines, an Artificial Intelligence Perspective*, RUN
- 2008-07 Peter van Rosmalen, *Supporting the Tutor in the Design and Support of Adaptive E-Learning*, OU
- 2008-08 Janneke Bolt, *Bayesian Networks: Aspects of Approximate Inference*, UU
- 2008-09 Christof van Nimwegen, *The Paradox of the Guided User: Assistance can be Counter-Effective*, UU
- 2008-10 Wauter Bosma, *Discourse Oriented Summarization*, UT
- 2008-11 Vera Kartseva, *Designing Controls for Network Organizations: A Value-Based Approach*, VU
- 2008-12 Jozsef Farkas, *A Semiotically Oriented Cognitive Model of Knowledge Representation*, RUN
- 2008-13 Caterina Carraciolo, *Topic Driven Access to Scientific Handbooks*, UvA
- 2008-14 Arthur van Bunningen, *Context-Aware Querying; Better Answers with Less Effort*, UT
- 2008-15 Martijn van Otterlo, *The Logic of Adaptive Behavior: Knowledge Representation and Algorithms for the Markov Decision Process Framework in First-Order Domains*, UT
- 2008-16 Henriette van Vugt, *Embodied Agents from a User's Perspective*, VU
- 2008-17 Martin Op 't Land, *Applying Architecture and Ontology to the Splitting and Allying of Enterprises*, TUD
- 2008-18 Guido de Croon, *Adaptive Active Vision*, UM
- 2008-19 Henning Rode, *From Document to Entity Retrieval: Improving Precision and Performance of Focused Text Search*, UT
- 2008-20 Rex Arendsen, *Geen Bericht, Goed Bericht. Een Onderzoek naar de Effecten van de Introductie van Elektronisch Berichtenverkeer met de Overheid op de Administratieve Lasten van Bedrijven*, UvA
- 2008-21 Krisztian Balog, *People Search in the Enterprise*, UvA
- 2008-22 Henk Koning, *Communication of IT-Architecture*, UU
- 2008-23 Stefan Visscher, *Bayesian Network Models for the Management of Ventilator-Associated Pneumonia*, UU
- 2008-24 Zharko Aleksovski, *Using Background Knowledge in Ontology Matching*, VU
- 2008-25 Geert Jonker, *Efficient and Equitable Exchange in Air Traffic Management Plan Repair Using Spender-Signed Currency*, UU
- 2008-26 Marijn Huijbregts, *Segmentation, Diarization and Speech Transcription: Surprise Data Unraveled*, UT
- 2008-27 Hubert Vogten, *Design and Implementation Strategies for IMS Learning Design*, OU
- 2008-28 Ildiko Flesch, *On the Use of Independence Relations in Bayesian Networks*, RUN
- 2008-29 Dennis Reidsma, *Annotations and Subjective Machines - Of Annotators, Embodied Agents, Users, and Other Humans*, UT
- 2008-30 Wouter van Atteveldt, *Semantic Network Analysis: Techniques for Extracting, Representing and Querying Media Content*, VU
- 2008-31 Loes Braun, *Pro-Active Medical Information Retrieval*, UM
- 2008-32 Trung H. Bui, *Toward Affective Dialogue Management using Partially Observable Markov Decision Processes*, UT
- 2008-33 Frank Terpstra, *Scientific Workflow Design; Theoretical and Practical Issues*, UvA
- 2008-34 Jeroen de Knijf, *Studies in Frequent Tree Mining*, UU
- 2008-35 Ben Torben Nielsen, *Dendritic Morphologies: Function Shapes Structure*, UvT

—2009—

2009-01 Rasa Jurgelenaite, *Symmetric Causal Independence Models*, RUN

2009-02 Willem Robert van Hage, *Evaluating Ontology-Alignment Techniques*, VU

2009-03 Hans Stol, *A Framework for Evidence-Based Policy Making Using IT*, UvT

2009-04 Josephine Nabukenya, *Improving the Quality of Organisational Policy Making Using Collaboration Engineering*, RUN

2009-05 Sietse Overbeek, *Bridging Supply and Demand for Knowledge Intensive Tasks - Based on Knowledge, Cognition, and Quality*, RUN

2009-06 Muhammad Subianto, *Understanding Classification*, UU

2009-07 Ronald Poppe, *Discriminative Vision-Based Recovery and Recognition of Human Motion*, UT

2009-08 Volker Nannen, *Evolutionary Agent-Based Policy Analysis in Dynamic Environments*, VU

2009-09 Benjamin Kanagwa, *Design, Discovery and Construction of Service-Oriented Systems*, RUN

2009-10 Jan Wielemaker, *Logic Programming for Knowledge-Intensive Interactive Applications*, UvA

2009-11 Alexander Boer, *Legal Theory, Sources of Law & the Semantic Web*, UvA

2009-12 Peter Massuthé, *Operating Guidelines for Services*, TUE/Humboldt-Universitaet zu Berlin

2009-13 Steven de Jong, *Fairness in Multi-Agent Systems*, UM

2009-14 Maksym Korotkiy, *From Ontology-Enabled Services to Service-Enabled Ontologies (Making Ontologies Work in E-Science with ONTO-SOA)*, VU

2009-15 Rinke Hoekstra, *Ontology Representation - Design Patterns and Ontologies that Make Sense*, UvA

2009-16 Fritz Reul, *New Architectures in Computer Chess*, UvT

2009-17 Laurens van der Maaten, *Feature Extraction from Visual Data*, UvT

2009-18 Fabian Groffen, *Armada, An Evolving Database System*, CWI

2009-19 Valentin Robu, *Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets*, CWI

2009-20 Bob van der Vecht, *Adjustable Autonomy: Controlling Influences on Decision Making*, UU

2009-21 Stijn Vanderlooy, *Ranking and Reliable Classification*, UM

2009-22 Pavel Serdyukov, *Search For Expertise: Going Beyond Direct Evidence*, UT

2009-23 Peter Hofgesang, *Modelling Web Usage in a Changing Environment*, VU

2009-24 Annerieke Heuvelink, *Cognitive Models for Training Simulations*, VU

2009-25 Alex van Ballegooij, *RAM: Array Database Management Through Relational Mapping*, CWI

2009-26 Fernando Koch, *An Agent-Based Model for the Development of Intelligent Mobile Services*, UU

2009-27 Christian Glahn, *Contextual Support of Social Engagement and Reflection on the Web*, OU

2009-28 Sander Evers, *Sensor Data Management with Probabilistic Models*, UT

2009-29 Stanislav Pokraev, *Model-Driven Semantic Integration of Service-Oriented Applications*, UT

2009-30 Marcin Zukowski, *Balancing Vectorized Query Execution with Bandwidth-Optimized Storage*, CWI

2009-31 Sofiya Katrenko, *A Closer Look at Learning Relations from Text*, UvA

2009-32 Rik Farenhorst, and Remco de Boer, *Architectural Knowledge Management: Supporting Architects and Auditors*, VU

2009-33 Khiet Truong, *How Does Real Affect Affect Affect Recognition In Speech?*, UT

2009-34 Inge van de Weerd, *Advancing in Software Product Management: An Incremental Method Engineering Approach*, UU

2009-35 Wouter Koelewijn, *Privacy en Politiegegevens; Over Geautomatiseerde Normatieve Informatie-Uitwisseling*, UL

2009-36 Marco Kalz, *Placement Support for Learners in Learning Networks*, OU

2009-37 Hendrik Drachslér, *Navigation Support for Learners in Informal Learning Networks*, OU

2009-38 Riina Vuorikari, *Tags and Self-Organisation: a Metadata E-cology for Learning Resources in a Multilingual Context*, OU

2009-39 Christian Stahl, *Service Substitution - A Behavioral Approach Based on Petri Nets*, TUE/Humboldt-Universitaet zu Berlin

2009-40 Stephan Raaijmakers, *Multinomial Language Learning: Investigations into the Geometry of Language*, UvT

2009-41 Igor Berezhnyy, *Digital Analysis of Paintings*, UvT

2009-42 Toine Bogers, *Recommender Systems for Social Bookmarking*, UvT

2009-43 Virginia Nunes Leal Franqueira, *Finding Multi-Step Attacks in Computer Networks using Heuristic Search and Mobile Ambients*, UT

2009-44 Roberto Santana Tapia, *Assessing Business-IT Alignment in Networked Organizations*, UT

2009-45 Jilles Vreeken, *Making Pattern Mining Useful*, UU

2009-46 Loredana Afanasiev, *Querying XML: Benchmarks and Recursion*, UvA

—2010—

2010-01 Matthijs van Leeuwen, *Patterns that Matter*, UU

2010-02 Ingo Wassink, *Work flows in Life Science*, UT

2010-03 Joost Geurts, *A Document Engineering Model and Processing Framework for Multimedia documents*, CWI

2010-04 Olga Kulyk, *Do You Know What I Know? Situational Awareness of Co-located Teams in Multidisplay Environments*, UT

2010-05 Claudia Hauff, *Predicting the Effectiveness of Queries and Retrieval Systems*, UT

2010-06 Sander Bakkes, *Rapid Adaptation of Video Game AI*, UvT

2010-07 Wim Fikkert, *Gesture Interaction at a Distance*, UT

2010-08 Krzysztof Siewicz, *Towards an Improved Regulatory Framework of Free Software. Protecting User Freedoms in a World of Software Communities and E-Governments*, UL

2010-09 Hugo Kielman, *Politiële Gegevensverwerking en Privacy, Naar een Effectieve Waarboring*, UL

- 2010-10** Rebecca Ong, *Mobile Communication and Protection of Children*, UL
- 2010-11** Adriaan Ter Mors, *The World According to MARP: Multi-Agent Route Planning*, TUD
- 2010-12** Susan van den Braak, *Sensemaking Software for Crime Analysis*, UU
- 2010-13** Gianluigi Folino, *High Performance Data Mining Using Bio-inspired techniques*, RUN
- 2010-14** Sander van Splunter, *Automated Web Service Reconfiguration*, VU
- 2010-15** Lianne Bodenstaff, *Managing Dependency Relations in Inter-Organizational Models*, UT
- 2010-16** Sicco Verwer, *Efficient Identification of Timed Automata, Theory and Practice*, TUD
- 2010-17** Spyros Kotoulas, *Scalable Discovery of Networked Resources: Algorithms, Infrastructure, Applications*, VU
- 2010-18** Charlotte Gerritsen, *Caught in the Act: Investigating Crime by Agent-Based Simulation*, VU
- 2010-19** Henriette Cramer, *People's Responses to Autonomous and Adaptive Systems*, UvA
- 2010-20** Ivo Swartjes, *Whose Story Is It Anyway? How Improv Informs Agency and Authorship of Emergent Narrative*, UT
- 2010-21** Harold van Heerde, *Privacy-Aware Data Management by Means of Data Degradation*, UT
- 2010-22** Michiel Hildebrand, *End-user Support for Access to Heterogeneous Linked Data*, CWI
- 2010-23** Bas Steunebrink, *The Logical Structure of Emotions*, UU
- 2010-24** Dmytro Tykhonov, *Designing Generic and Efficient Negotiation Strategies*
- 2010-25** Zulfiqar Ali Memon, *Modelling Human-Awareness for Ambient Agents: A Human Mindreading Perspective*, VU
- 2010-26** Ying Zhang, *XRPC: Efficient Distributed Query Processing on Heterogeneous XQuery Engines*, CWI
- 2010-27** Marten Voulon, *Automatisch Contracteren*, UL
- 2010-28** Arne Koopman, *Characteristic Relational Patterns*, UU
- 2010-29** Stratos Idreos, *Database Cracking: Towards Auto-Tuning Database Kernels*, CWI
- 2010-30** Marieke van Erp, *Accessing Natural History - Discoveries in Data Cleaning, Structuring, and Retrieval*, UvT
- 2010-31** Victor de Boer, *Ontology Enrichment from Heterogeneous Sources on the Web*, UvA
- 2010-32** Marcel Hiel, *An Adaptive Service Oriented Architecture: Automatically solving Interoperability Problems*, UvT
- 2010-33** Robin Aly, *Modeling Representation Uncertainty in Concept-Based Multimedia Retrieval*, UT
- 2010-34** Teduh Dirgahayu, *Interaction Design in Service Compositions*, UT
- 2010-35** Dolf Trieschnigg, *Proof of Concept: Concept-based Biomedical Information Retrieval*, UT
- 2010-36** Jose Janssen, *Paving the Way for Lifelong Learning: Facilitating Competence Development Through a Learning Path Specification*, OU
- 2010-37** Niels Lohmann, *Correctness of Services and Their Composition*, TUE
- 2010-38** Dirk Fahland, *From Scenarios to Components*, TUE
- 2010-39** Ghazanfar Farooq Siddiqui, *Integrative Modeling of Emotions in Virtual Agents*, VU
- 2010-40** Mark van Assem, *Converting and Integrating Vocabularies for the Semantic Web*, VU
- 2010-41** Guillaume Chaslot, *Monte-Carlo Tree Search*, UM
- 2010-42** Sybren de Kinderen, *Needs-Driven Service Bundling in a Multi-Supplier Setting - the Computational E3-Service Approach*, VU
- 2010-43** Peter van Kranenburg, *A Computational Approach to Content-Based Retrieval of Folk Song Melodies*, UU
- 2010-44** Pieter Bellekens, *An Approach Towards Context-Sensitive and User-adapted Access to Heterogeneous Data Sources, Illustrated in the Television Domain*, TUE
- 2010-45** Vasilios Andrikopoulos, *A Theory and Model for the Evolution of Software Services*, UvT
- 2010-46** Vincent Pijpers, *e3alignment: Exploring Inter-Organizational Business-ICT Alignment*, VU
- 2010-47** Chen Li, *Mining Process Model Variants: Challenges, Techniques, Examples*, UT
- 2010-48** Milan Lovric, *Behavioral Finance and Agent-Based Artificial Markets*, EUR
- 2010-49** Jahn-Takeshi Saito, *Solving Difficult Game Positions*, UM
- 2010-50** Bouke Huurnink, *Search in Audiovisual Broadcast Archives*, UvA
- 2010-51** Alia Khairia Amin, *Understanding and Supporting Information Seeking Tasks in Multiple Sources*, CWI
- 2010-52** Peter-Paul van Maanen, *Adaptive Support for Human-Computer Teams: Exploring the Use of Cognitive Models of Trust and Attention*, VU
- 2010-53** Edgar Meij, *Combining Concepts and Language Models for Information Access*, UvA
- 2011—
- 2011-01** Botond Cseke, *Variational Algorithms for Bayesian Inference in Latent Gaussian Models*, RUN
- 2011-02** Nick Tinnemeier, *Work Flows in Life Science*, UU
- 2011-03** Jan Martijn van der Werf, *Compositional Design and Verification of Component-Based Information Systems*, TUE
- 2011-04** Hado van Hasselt, *Insights in Reinforcement Learning: Formal Analysis and Empirical Evaluation of Temporal-Difference Learning Algorithms*, UU
- 2011-05** Base van der Raadt, *Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline*, VU
- 2011-06** Yiwen Wang, *Semantically-Enhanced Recommendations in Cultural Heritage*, TUE
- 2011-07** Yujia Cao, *Multimodal Information Presentation for High Load Human Computer Interaction*, UT
- 2011-08** Nieske Vergunst, *BDI-Based Generation of Robust Task-Oriented Dialogues*, UU
- 2011-09** Tim de Jong, *Contextualised Mobile Media for Learning*, OU
- 2011-10** Bart Bogaert, *Cloud Content Contention*, UvT
- 2011-11** Dhaval Vyas, *Designing for Awareness: An Experience-Focused HCI Perspective*, UT
- 2011-12** Carmen Bratosin, *Grid Architecture for Distributed Process Mining*, TUE
- 2011-13** Xiaoyu Mao, *Airport Under Control. Multi-Agent Scheduling for Airport Ground Handling*, UvT
- 2011-14** Milan Lovric, *Behavioral Finance and Agent-Based Artificial Markets*, EUR

- 2011-15** Marijn Koolen, *The Meaning of Structure: the Value of Link Evidence for Information Retrieval*, UvA
- 2011-16** Maarten Schadd, *Selective Search in Games of Different Complexity*, UM
- 2011-17** Jiyin He, *Exploring Topic Structure: Coherence, Diversity and Relatedness*, UvA
- 2011-18** Mark Ponsen, *Strategic Decision-Making in Complex Games*, UM
- 2011-19** Ellen Rusman, *The Mind 's Eye on Personal Profiles*, OU
- 2011-20** Qing Gu, *Guiding Service-Oriented Software Engineering - A View-Based Approach*, VU
- 2011-21** Linda Terlouw, *Modularization and Specification of Service-Oriented Systems*, TUD
- 2011-22** Junte Zhang, *System Evaluation of Archival Description and Access*, UvA
- 2011-23** Wouter Weerkamp, *Finding People and their Utterances in Social Media*, UvA
- 2011-24** Herwin van Welbergen, *Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior*, UT
- 2011-25** Syed Waqar ul Qounain Jaffry, *Analysis and Validation of Models for Trust Dynamics*, VU
- 2011-26** Matthijs Aart Pontier, *Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots*, VU
- 2011-27** Aniel Bhulai, *Dynamic Website Optimization Through Autonomous Management of Design Patterns*, VU
- 2011-28** Rianne Kaptein, *Effective Focused Retrieval by Exploiting Query Context and Document Structure*, UvA
- 2011-29** Faisal Kamiran, *Discrimination-Aware Classification*, TUE
- 2011-30** Egon van den Broek, *Affective Signal Processing (ASP): Unraveling the Mystery of Emotions*, UT
- 2011-31** Ludo Waltman, *Computational and Game-Theoretic Approaches for Modeling Bounded Rationality*, EUR
- 2011-32** Nees-Jan van Eck, *Methodological Advances in Bibliometric Mapping of Science*, EUR
- 2011-33** Tom van der Weide, *Arguing to Motivate Decisions*, UU
- 2011-34** Paolo Turrini, *Strategic Reasoning in Interdependence: Logical and Game-Theoretical Investigations*, UU
- 2011-35** Maaïke Harbers, *Explaining Agent Behavior in Virtual Training*, UU
- 2011-36** Erik van der Spek, *Experiments in Serious Game Design: a Cognitive Approach*, UU
- 2011-37** Adriana Burlutiu, *Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference*, RUN
- 2011-38** Nyree Lemmens, *Bee-inspired Distributed Optimization*, UM
- 2011-39** Joost Westra, *Organizing Adaptation Using Agents in Serious Games*, UU
- 2011-40** Viktor Clerc, *Architectural Knowledge Management in Global Software Development*, VU
- 2011-41** Luan Ibraimi, *Cryptographically Enforced Distributed Data Access Control*, UT
- 2011-42** Michal Sindlar, *Explaining Behavior Through Mental State Attribution*, UU
- 2011-43** Henk van der Schuur, *Process Improvement Through Software Operation Knowledge*, UU
- 2011-44** Boris Reuderink, *Robust Brain-Computer Interfaces*, UT
- 2011-45** Herman Stehouwer, *Statistical Language Models for Alternative Sequence Selection*, UvT
- 2011-46** Beibei Hu, *Towards Contextualized Information Delivery: A Rule-Based Architecture for the Domain of Mobile Police Work*, TUD
- 2011-47** Azizi Bin Ab Aziz, *Exploring Computational Models for Intelligent Support of Persons with Depression*, VU
- 2011-48** Mark Ter Maat, *Response Selection and Turn-Taking for a Sensitive Artificial Listening Agent*, UT
- 2011-49** Andreea Niculescu, *Conversational Interfaces for Task-Oriented Spoken Dialogues: Design Aspects Influencing Interaction Quality*, UT
- 2012—
- 2012-01** Terry Kakeeto, *Relationship Marketing for SMEs in Uganda*, UvT
- 2012-02** Muhammad Umair, *Adaptivity, Emotion, and Rationality in Human and Ambient Agent Models*, VU
- 2012-03** Adam Vanya, *Supporting Architecture Evolution by Mining Software Repositories*, VU
- 2012-04** Jurriaan Souer, *Development of Content Management System-based Web Applications*, UU
- 2012-05** Marijn Plomp, *Maturing Inter-Organisational Information Systems*, UU
- 2012-06** Wolfgang Reinhardt, *Awareness Support for Knowledge Workers in Research Networks*, OU
- 2012-07** Rianne van Lambalgen, *When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions*, VU
- 2012-08** Gerben de Vries, *Kernel Methods for Vessel Trajectories*, UvA
- 2012-09** Ricardo Neisse, *Trust and Privacy Management Support for Context-Aware Service Platforms*, UT
- 2012-10** David Smits, *Towards a Generic Distributed Adaptive Hypermedia Environment*, TUE
- 2012-11** J.C.B. Rantham Prabhakara, *Process Mining in the Large: Preprocessing, Discovery, and Diagnostics*, TUE
- 2012-12** Kees van der Sluijs, *Model Driven Design and Data Integration in Semantic Web Information Systems*, TUE
- 2012-13** Suleman Shahid, *Fun and Face: Exploring Non-Verbal Expressions of Emotion During Playful Interactions*, UvT
- 2012-14** Evgeny Knutov(TUE) *Generic Adaptation Framework for Unifying Adaptive Web-based Systems*, TUE
- 2012-15** Natalie van der Wal, *Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes*, VU
- 2012-16** Fiemke Both, *Helping People by Understanding Them - Ambient Agents Supporting Task Execution and Depression Treatment*, VU
- 2012-17** Amal Elgammal, *Towards a Comprehensive Framework for Business Process Compliance*, UvT
- 2012-18** Eltjo Poort, *Improving Solution Architecting Practices*, VU

- 2012-19** Helen Schonenberg, *What's Next? Operational Support for Business Process Execution*, TUE
- 2012-20** Ali Bahramisharif, *Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing*, RUN
- 2012-21** Roberto Cornacchia, *Querying Sparse Matrices for Information Retrieval*, TUD
- 2012-22** Thijs Vis, *Intelligence, Politie en Veiligheidsdienst: Verenigbare Grootheden?*, UvT
- 2012-23** Christian Muehl, *Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction*, UT
- 2012-24** Laurens van der Werff, *Evaluation of Noisy Transcripts for Spoken Document Retrieval*, UT
- 2012-25** Silja Eckartz, *Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application*, UT
- 2012-26** Emile de Maat, *Making Sense of Legal Text*, UvA
- 2012-27** Hayrettin Gurkok, *Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games*, UT
- 2012-28** Nancy Pascall, *Engendering Technology Empowering Women*, UvT
- 2012-29** Almer Tigelaar, *Peer-to-Peer Information Retrieval*, UT
- 2012-30** Alina Pommeranz, *Designing Human-Centered Systems for Reflective Decision Making*, TUD
- 2012-31** Emily Bagarukayo, *A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure*, RUN
- 2012-32** Wietske Visser, *Qualitative Multi-Criteria Preference Representation and Reasoning*, TUD
- 2012-33** Rory Sie, *Coalitions in Cooperation Networks (CO-COON)*, OU
- 2012-34** Pavol Jancura, *Evolutionary Analysis in PPI Networks and Applications*, RUN
- 2012-35** Evert Haasdijk, *Never Too Old To Learn – On-line Evolution of Controllers in Swarm- and Modular Robotics*, VU
- 2012-36** Denis Ssebugwawo, *Analysis and Evaluation of Collaborative Modeling Processes*, RUN
- 2012-37** Agnes Nakakawa, *A Collaboration Process for Enterprise Architecture Creation*, RUN
- 2012-38** Selmar Smit, *Parameter Tuning and Scientific Testing in Evolutionary Algorithms*, VU
- 2012-39** Hassan Fatemi, *Risk-Aware Design of Value and Coordination Networks*, UT
- 2012-40** Agus Gunawan, *Information Access for SMEs in Indonesia*, UvT
- 2012-41** Sebastian Kelle, *Game Design Patterns for Learning*, OU
- 2012-42** Dominique Verpoorten, *Reflection Amplifiers in Self-Regulated Learning*, OU
- 2012-43** *Withdrawn*
- 2012-44** Anna Tordai, *On Combining Alignment Techniques*, VU
- 2012-45** Benedikt Kratz, *A Model and Language for Business-aware Transactions*, UvT
- 2012-46** Simon Carter, *Exploration and Exploitation of Multilingual Data for Statistical Machine Translation*, UvA
- 2012-47** Manos Tsagkias, *Mining Social Media: Tracking Content and Predicting Behavior*, UvA
- 2012-48** Jorn Bakker, *Handling Abrupt Changes in Evolving Time-series Data*, TUE
- 2012-49** Michael Kaisers, *Learning Against Learning - Evolutionary Dynamics of Reinforcement Learning Algorithms in Strategic Interactions*, UM
- 2012-50** Steven van Kervel, *Ontology-driven Enterprise Information Systems Engineering*, TUD
- 2012-51** Jeroen de Jong, *Heuristics in Dynamic Sceduling; a Practical Framework with a Case Study in Elevator Dispatching*, TUD
- 2013—
- 2013-01** Viorel Milea, *News Analytics for Financial Decision Support*, EUR
- 2013-02** Erietta Liarou, *MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing*, CWI
- 2013-03** Szymon Klarman, *Reasoning with Contexts in Description Logics*, VU
- 2013-04** Chetan Yadati, *Coordinating Autonomous Planning and Scheduling*, TUD
- 2013-05** Dulce Pumareja, *Groupware Requirements Evolutions Patterns*, UT
- 2013-06** Romulo Goncalves, *The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience*, CWI
- 2013-07** Giel van Lankveld, *Quantifying Individual Player Differences*, UT
- 2013-08** Robbert-Jan Merk, *Making Enemies: Cognitive Modeling for Opponent Agents in Fighter Pilot Simulators*, VU
- 2013-09** Fabio Gori, *Metagenomic Data Analysis: Computational Methods and Applications*, RUN
- 2013-10** Jeewanie Jayasinghe Arachchige, *A Unified Modeling Framework for Service Design*, UvT
- 2013-11** Evangelos Pournaras, *Multi-level Reconfigurable Self-organization in Overlay Services*, TUD
- 2013-12** Marian Razavian, *Knowledge-driven Migration to Services*, VU
- 2013-13** Mohammad Safiri, *Service Tailoring: User-centric Creation of Integrated IT-based Homecare Services to Support Independent Living of Elderly*, UT
- 2013-14** Jafar Tanha, *Ensemble Approaches to Semi-Supervised Learning*, UvA
- 2013-15** Daniel Hennes, *Multiagent Learning - Dynamic Games and Applications*, UM
- 2013-16** Eric Kok, *Exploring the Practical Benefits of Argumentation in Multi-agent Deliberation*, UU
- 2013-17** Koen Kok, *The PowerMatcher: Smart Coordination for the Smart Electricity Grid*, VU
- 2013-18** Jeroen Janssens, *Outlier Selection and One-Class Classification*, UvT
- 2013-19** Renze Steenhuisen, *Coordinated Multi-Agent Planning and Scheduling*, TUD
- 2013-20** Katja Hofmann, *Fast and Reliable Online Learning to Rank for Information Retrieval*, UvA
- 2013-21** Sander Wubben, *Text-to-text Generation by Monolingual Machine Translation*, UvT
- 2013-22** Tom Claassen, *Causal Discovery and Logic*, RUN
- 2013-23** Patricio de Alencar Silva, *Value Activity Monitoring*, UvT

- 2013-24 Haitham Bou Ammar, *Automated Transfer in Reinforcement Learning*, UM
- 2013-25 Agnieszka Anna Latoszek-Berendsen, *Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System*, UM
- 2013-26 Alireza Zarghami, *Architectural Support for Dynamic Homecare Service Provisioning*, UT
- 2013-27 Mohammad Huq, *Inference-based Framework Managing Data Provenance*, UT
- 2013-28 Frans van der Sluis, *When Complexity Becomes Interesting: An Inquiry into the Information eXperience*, UT
- 2013-29 Iwan de Kok, *Listening Heads*, UT
- 2013-30 Joyce Nakatumba, *Resource-Aware Business Process Management: Analysis and Support*, TUE
- 2013-31 Dinh Khoa Nguyen, *Blueprint Model and Language for Engineering Cloud Applications*, UvT
- 2013-32 Kamakshi Rajagopal, *Networking For Learning: The role of Networking in a Lifelong Learner's Professional Development*, OUN
- 2013-33 Qi Gao, *User Modeling and Personalization in the Microblogging Sphere*, TUD
- 2013-34 Kien Tjin-Kam-Jet (UT), *Distributed Deep Web Search*, UT
- 2013-35 Abdallah El Ali, *Minimal Mobile Human Computer Interaction*, UVA
- 2013-36 Than Lam Hoang, *Pattern Mining in Data Streams*, TUE
- 2013-37 Dirk Börner, *Ambient Learning Displays*, OUN
- 2013-38 Eelco den Heijer, *Autonomous Evolutionary Art*, VU
- 2013-39 Joop de Jong, *A Method for Enterprise Ontology based Design of Enterprise Information Systems*, TUD
- 2013-40 Pim Nijssen, *Monte-Carlo Tree Search for Multi-Player Games*, UM
- 2013-41 Jochem Liem, *Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning*, UVA
- 2013-42 Léon Planken, *Algorithms for Simple Temporal Reasoning*, TUD
- 2013-43 Marc Bron, *Exploration and Contextualization through Interaction and Concepts*, UVA
- 2014—
- 2014-01 Nicola Barile, *Studies in Learning Monotone Models from Data*, UU
- 2014-02 Fiona Tuliayo, *Combining System Dynamics with a Domain Modeling Method*, RUN
- 2014-03 Sergio Raul Duarte Torres, *Information Retrieval for Children: Search Behavior and Solutions*, UT
- 2014-04 Hanna Jochmann-Mannak, *Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation*, UT
- 2014-05 Jurriaan van Reijssen, *Knowledge Perspectives on Advancing Dynamic Capability*, UU
- 2014-06 Damian Tamburri, *Supporting Networked Software Development*, VU
- 2014-07 Arya Adriansyah, *Aligning Observed and Modeled Behavior*, TUE
- 2014-08 Samur Araujo, *Data Integration over Distributed and Heterogeneous Data Endpoints*, TUD
- 2014-09 Philip Jackson, *Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language*, UvT
- 2014-10 Ivan Salvador Razo Zapata, *Service Value Networks*, VU
- 2014-11 Janneke van der Zwaan, *An Empathic Virtual Buddy for Social Support*, TUD
- 2014-12 Willem van Willigen, *Look Ma, No Hands: Aspects of Autonomous Vehicle Control*, VU
- 2014-13 Arlette van Wissen, *Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains*, VU
- 2014-14 Yangyang Shi, *Language Models With Meta-information*, TUD
- 2014-15 Natalya Mogles, *Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare*, VU
- 2014-16 Krystyna Milián, *Supporting trial recruitment and design by automatically interpreting eligibility criteria*, VU
- 2014-17 Kathrin Dentler VU, *Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability*, VU
- 2014-18 Mattijs Ghijsen, *Methods and Models for the Design and Study of Dynamic Agent Organizations*, UVA
- 2014-19 Vinicius Ramos, *Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support*, TUE
- 2014-20 Mena Habib, *Named Entity Extraction and Disambiguation for Informal Text: The Missing Link*, UT
- 2014-21 Cassidy Clark, *Negotiation and Monitoring in Open Environments*, TUD
- 2014-22 Marieke Peeters, *Personalized Educational Games - Developing agent-supported scenario-based training*, UU
- 2014-23 Eleftherios Sidirourgos, *Space Efficient Indexes for the Big Data Era*, UvA/CWI
- 2014-24 Davide Ceolin, *Trusting Semi-structured Web Data*, VU
- 2014-25 2014-26 Tim Baarslag, *What to Bid and When to Stop*, TUD
- 2014-27 Rui Jorge Almeida, *Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty*, EUR
- 2014-28 Anna Chmielowiec, *Decentralized k-Clique Matching*, VU
- 2014-29 Jaap Kabbedijk, *Variability in Multi-Tenant Enterprise Software*, UU
- 2014-30 Peter de Cock, *Anticipating Criminal Behaviour*, UvT
- 2014-31 Leo van Moergestel, *Agent Technology in Agile Multi-parallel Manufacturing and Product Support*, UU
- 2014-32 Naser Ayat, *On Entity Resolution in Probabilistic Data*, UvA
- 2014-33 Tesfa Tegegne, *Service Discovery in eHealth*, RUN
- 2014-34 Christina Manteli, *The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems*, VU
- 2014-35 Joost van Ooijen, *Cognitive Agents in Virtual Worlds: A Middleware Design Approach*, UU
- 2014-36 Joos Buijs, *Flexible Evolutionary Algorithms for Mining Structured Process Models*, TUE
- 2014-37 Maral Dadvar, *Experts and Machines United Against Cyberbullying*, UT
- 2014-38 Danny Plass-Oude Bos, *Making brain-computer interfaces better: improving usability through post-processing*, UT
- 2014-39 Jasmína Maric, *Web Communities, Immigration, and Social Capital*, UvT
- 2014-40 Walter Omona, *A Framework for Knowledge Management Using ICT in Higher Education*, RUN
- 2014-41 Frederic Hogenboom, *Automated Detection of Financial Events in News Text*, EUR
- 2014-42 Carsten Eijckhof, *Contextual Multidimensional Relevance Models*, CWI/TUD
- 2014-43 Kevin Vlaanderen, *Supporting Process Improvement*

using *Method Increments*, UU

2014-44 Paulien Meesters, *Intelligent Blauw. Met als ondertitel: Intelligentie-gestuurde politiezorg in gebiedsgebonden eenheden.*, UvT

2014-45 Birgit Schmitz, *Mobile Games for Learning: A Pattern-Based Approach*, OUN

2014-46 Ke Tao, *Social Web Data Analytics: Relevance, Redundancy, Diversity*, TUD

2014-47 Shangsong Liang, *Fusion and Diversification in Information Retrieval*, UVA

—2015—

2015-01 Niels Netten, *Machine Learning for Relevance of Information in Crisis Response*, UvA

2015-02 Faiza Bukhsh, *Smart auditing: Innovative Compliance Checking in Customs Controls*, UvT

2015-03 Twan van Laarhoven, *Machine learning for network data*, RUN

2015-04 Howard Spoelstra (OUN), *Collaborations in Open Learning Environments*, OUN

2015-05 Christoph Bösch, *Cryptographically Enforced Search Pattern Hiding*, UT

2015-06 Farideh Heidari, *Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes*, TUD

2015-07 Maria-Hendrike Peetz, *Time-Aware Online Reputation Analysis*, UvA

2015-08 Jie Jiang, *Organizational Compliance: An agent-based model for designing and evaluating organizational interactions*, TUD

2015-09 Randy Klaassen, *HCI Perspectives on Behavior Change Support Systems*, UT

2015-10 Henry Hermans, *OpenU: design of an integrated system to support lifelong learning*, OUN

2015-11 Yongming Luo, *Designing algorithms for big graph datasets: A study of computing bisimulation and joins*, TUE

2015-12 Julie M. Birkholz, *Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks*, VU

2015-13 Giuseppe Procaccianti (VU), *Energy-Efficient Software*, VU

2015-14 Bart van Straalen, *A cognitive approach to modeling bad news conversations*, UT

2015-15 Klaas Andries de Graaf, *Ontology-based Software Architecture Documentation*, VU

2015-16 Changyun Wei, *Cognitive Coordination for Cooperative Multi-Robot Teamwork*, UT

2015-17 André van Cleeff, *Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs*, UT

2015-18 Holger Pirk, *Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories*, CWI

2015-19 Bernardo Tabuenca (OUN), *Ubiquitous Technology for Lifelong Learners*, OUN

2015-20 Lois Vanhée, *Using Culture and Values to Support Flexible Coordination*, UU

2015-21 Sibren Fetter (OUN), *Using Peer-Support to Expand and Stabilize Online Learning*, OUN

2015-22 **2015-23** Luit Gazendam, *Cataloguer Support in Cultural Heritage*, VU

2015-24 Richard Berendsen, *Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation*, UVA

2015-25 Steven Woudenberg, *Bayesian Tools for Early Disease Detection*, UU

2015-26 Alexander Hogenboom, *Sentiment Analysis of Text Guided by Semantics and Structure*, EUR

2015-27 Sándor Héman, *Updating compressed column stores*,

CWI

2015-28 Janet Bagorogoza, *Knowledge Management and High Performance; The Uganda Financial Institutions Model for HPO*, TiU

2015-29 Hendrik Baier, *Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains*, UM

2015-30 Kiavash Bahreini, *Real-time Multimodal Emotion Recognition in E-Learning*, OU

2015-31 Yakup Koç, *On the robustness of Power Grids*, TUD

2015-32 Jerome Gard, *Corporate Venture Management in SMEs*, UL

2015-33 Frederik Schadd, *Ontology Mapping with Auxiliary Resources*, TUD

2015-34 Victor de Graaf, *Gesocial Recommender Systems*, UT

2015-35 Jungxiao Xu, *Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction*, TUD

—2016—

2016-01 Syed Saiden Abbas, *Recognition of Shapes by Humans and Machines*, RUN

2016-02 Michiel Christiaan Meulendijk, *Optimizing medication reviews through decision support: prescribing a better pill to swallow*, UU

2016-03 Maya Sappelli, *Knowledge Work in Context: User Centered Knowledge Worker Support*, RUN

2016-04 Laurens Rietveld, *Publishing and Consuming Linked Data*, VU

2016-05 Evgeny Sherkhonov, *Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers*, UVA

2016-06 Michel Wilson, *Robust scheduling in an uncertain environment*, TUD

2016-07 Jeroen de Man, *Measuring and modeling negative emotions for virtual training*, VU

2016-08 Matje van de Camp (TiU), *A Link to the Past: Constructing Historical Social Networks from Unstructured Data*, TiU

2016-09 Archana Nottamkandath, *Trusting Crowdsourced Information on Cultural Artefacts*, VU

2016-10 George Karafotias, *Parameter Control for Evolutionary Algorithms*, VUA

2016-11 Anne Schuth, *Search Engines that Learn from Their Users*, UVA

2016-12 Max Knobbout, *Logics for Modelling and Verifying Normative Multi-Agent Systems*, UU

2016-13 Nana Baah Gyan, *The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach*, VU

2016-14 Ravi Khadka, *Revisiting Legacy Software System Modernization*, UU

2016-15 Steffen Michels, *Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments*, RUN

2016-16 Guangliang Li, *Socially Intelligent Autonomous Agents that Learn from Human Reward*, UVA

2016-17 Berend Weel, *Towards Embodied Evolution of Robot Organisms*, VU

2016-18 Albert Meroño Peñuela, *Refining Statistical Data on the Web*, VU

2016-19 Julia Efreмова, *Mining Social Structures from Genealogical Data*, TU/e

2016-20 Daan Odijk, *Context & Semantics in News & Web Search*, UVA

2016-21 Alejandro Moreno Cèlleri, *From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground*, UT

2016-22 Grace Lewis, *Software Architecture Strategies for Cyber-Foraging Systems*, VU

- 2016-23 Fei Cai, *Query Auto Completion in Information Retrieval*, UVA
- 2016-24 Brend Wanders, *Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach*, UT
- 2016-25 Julia Kiseleva, *Using Contextual Information to Understand Searching and Browsing Behavior*, TU/e
- 2016-26 Dilhan Thilakarathne, *In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains*, VU
- 2016-27 Wen Li, *Understanding Geo-spatial Information on Social Media*, TUD
- 2016-28 Mingxin Zhang, *Large-scale Agent-based Social Simulation - A study on epidemic prediction and control*, TUD
- 2016-29 Nicolas Höning, *Peak reduction in decentralised electricity systems - Markets and prices for flexible planning*, TUD
- 2016-30 Ruud Mattheij, *The Eyes Have It*, UvT
- 2016-31 Mohammad Khelghati (UT), *Deep web content monitoring*, UT
- 2016-32 Eelco Vriezepak, *Assessing Telecommunication Service Availability Risks for Crisis Organisations*, UT
- 2016-33 Peter Bloem, *Single Sample Statistics, exercises in learning from just one example*, UVA
- 2016-34 Dennis Schunselaar, *Configurable Process Trees: Elicitation, Analysis, and Enactment*, TUE
- 2016-35 Zhaochun Ren, *Monitoring Social Media: Summarization, Classification and Recommendation*, UVA
- 2016-36 Daphne Karreman, *Beyond R2D2: The design of non-verbal interaction behavior optimized for robot-specific morphologies*, UT
- 2016-37 Giovanni Sileno (UvA), *Aligning Law and Action - a conceptual and computational inquiry*, UvA
- 2016-38 Andrea Minuto, *Materials that Matter - Smart Materials meet Art & Interaction Design*, UT
- 2016-39 Merijn Bruijnes, *Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect*, UT
- 2016-40 Christian Detweiler, *Accounting for Values in Design*, TUD
- 2016-41 Thomas King, *Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance*, TUD
- 2016-42 Spyros Martzoukos, *Combinatorial and Compositional Aspects of Bilingual Aligned Corpora*, UVA
- 2016-43 Saskia Koldijk, *Context-Aware Support for Stress Self-Management: From Theory to Practice*, RUN
- 2016-44 Thibault Sellam, *Automatic Assistants for Database Exploration*, UVA
- 2016-45 Bram van de Laar, *Experiencing Brain-Computer Interface Control*, UT
- 2016-46 Jorge Gallego Perez, *Robots to Make you Happy*, UT
- 2016-47 Christina Weber, *Real-time foresight - Preparedness for dynamic innovation networks*, UL
- 2016-48 Tanja Buttler, *Collecting Lessons Learned*, TUD
- 2016-49 Gleb Polevoy, *Participation and Interaction in Projects. A Game-Theoretic Analysis*, TUD
- 2016-50 Yan Wang, *The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains*, UvT
- 2017-01 Jan-Jaap Oerlemans, *Investigating Cybercrime*, UL
- 2017-02 Sjoerd Timmer, *Designing and Understanding Forensic Bayesian Networks using Argumentation*, UU
- 2017-03 Daniël Harold Telgen, *Grid Manufacturing: A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines*, UU
- 2017-04 Mrunal Gawade, *Multi-core Parallelism in a Column-store*, CWI
- 2017-05 Mahdieh Shadi, *Collaboration Behavior*, UVA
- 2017-06 Damir Vandić, *Intelligent Information Systems for Web Product Search*, EUR
- 2017-07 Roel Bertens, *Insight in Information: from Abstract to Anomaly*, UU
- 2017-08 Rob Konijn, *Detecting Interesting Differences: Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery*, VU
- 2017-09 Dong Nguyen, *Text as Social and Cultural Data: A Computational Perspective on Variation in Text*, UT
- 2017-10 Robby van Delden, *(Steering) Interactive Play Behavior*, UT
- 2017-11 Florian Kunneman, *Modelling patterns of time and emotion in Twitter #anticipointment*, RUN
- 2017-12 Sander Leemans, *Robust Process Mining with Guarantees*, TUE
- 2017-13 Gijs Huisman, *Social Touch Technology - Extending the reach of social touch through haptic technology*, UT
- 2017-14 Shoshannah Tekofsky, *You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior*, UvT
- 2017-15 Peter Berck, *Memory-Based Text Correction*, RUN
- 2017-16 Aleksandr Chuklin, *Understanding and Modeling Users of Modern Search Engines*, UVA
- 2017-17 Daniel Dimov, *Crowdsourced Online Dispute Resolution*, UL
- 2017-18 Ridho Reinanda, *Entity Associations for Search*, UVA
- 2017-19 Jeroen Vuurens, *Proximity of Terms, Texts and Semantic Vectors in Information Retrieval*, UT
- 2017-20 Mohammadbashir Sedighi, *Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility*, TUD
- 2017-21 Jeroen Linssen, *Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)*, UT
- 2017-22 Sara Magliacane, *Logics for causal inference under uncertainty*, VU
- 2017-23 David Graus, *Entities of Interest — Discovery in Digital Traces*, UVA
- 2017-24 Chang Wang (TUD), *Use of Affordances for Efficient Robot Learning*, TUD
- 2017-25 Veruska Zamborlini, *Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search*, VU
- 2017-29 Merel Jung, *Socially intelligent robots that understand and respond to human touch*, UT
- 2017-27 Michiel Joosse, *Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors*, UT
- 2017-28 John Klein, *Architecture Practices for Complex Contexts*, VU
- 2017-29 Adel Alhuraibi, *From IT-Business Strategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT*, UvT
- 2017-30 Wilma Latuny, *The Power of Facial Expressions*, UvT
- 2017-31 Ben Ruijl, *Advances in computational methods for QFT calculations*, UL
- 2017-32 Thaeer Samar, *Access to and Retrieval of Content in Web Archives*, RUN
- 2017-33 Brigit van Loggem, *Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity*, OU
- 2017-34 Maren Scheffel, *The Evaluation Framework for Learning Analytics*, OU
- 2017-35 Martine de Vos, *Interpreting natural science spreadsheets*, VU
- 2017-36 Yuanhao Guo, *Shape Analysis for Phenotype Charac-*

terisation from High-throughput Imaging, UL

2017-37 Alejandro Montes Garcia , *WiBAF: A Within Browser*

Adaptation Framework that Enables Control over Privacy, TUE

2017-38 Alex Kayal, *Normative Social Applications*, TUD

2017-39 Sara Ahmadi, *Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR*, RUN

2017-40 Altaf Hussain Abro, *Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems*, VUA

2017-41 Adnan Manzoor, *Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle*, VUA

2017-42 Elena Sokolova, *Causal discovery from mixed and missing data with applications on ADHD datasets*, RUN

2017-43 Maaïke de Boer, *Semantic Mapping in Video Retrieval*, RUN

2017-44 Garm Lucassen, *Understanding User Stories - Computational Linguistics in Agile Requirements Engineering*, UU

2017-45 Bas Testerink, *Decentralized Runtime Norm Enforcement*, UU

2017-46 Jan Schneider, *Sensor-based Learning Support*, OU

2017-47 Jie Yang, *Crowd Knowledge Creation Acceleration*, TUD

2017-48 Angel Suarez, *Collaborative inquiry-based learning*, OU

—2018—

2018-01 Han van der Aa, *Comparing and Aligning Process Representations*, VUA

2018-02 Felix Mannhardt, *Multi-perspective Process Mining*, TUE

2018-03 Steven Bosems, *Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction*, UT

2018-04 Jordan Janeiro, *Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks*, TUD

2018-05 Hugo Huurdeman, *Supporting the Complex Dynamics of the Information Seeking Process*, UVA

2018-06 Dan Ionita, *Model-Driven Information Security Risk Assessment of Socio-Technical Systems*, UT

2018-07 Jieting Luo, *A formal account of opportunism in multi-agent systems*, UU

2018-08 Rick Smetsers, *Advances in Model Learning for Software Systems*, RUN

2018-09 Xu Xie, *Data Assimilation in Discrete Event Simulations*, TUD

2018-10 Julienka Mollee, *Moving forward: supporting physical activity behavior change through intelligent technology*, VUA

2018-11 Mahdi Sargolzaei, *Enabling Framework for Service-oriented Collaborative Networks*, UVA

2018-12 Xixi Lu, *Using behavioral context in process mining*, TUE

2018-13 Seyed Amin Tabatabaei, *Computing a Sustainable Future*, VUA

2018-14 Bart Joosten, *Detecting Social Signals with Spatiotemporal Gabor Filters*, UVT

2018-15 Naser Davarzani, *Biomarker discovery in heart failure*, UM

2018-16 Jaebok Kim, *Automatic recognition of engagement and emotion in a group of children*, UT

2018-17 Jianpeng Zhang, *On Graph Sample Clustering*, TUE

2018-18 Henriette Nakad, *De Notaris en Private Rechtspraak*, UL

2018-19 Minh Duc Pham, *Emergent relational schemas for RDF*,

VUA

2018-20 Manxia Liu, *Time and Bayesian Networks*, RUN

2018-21 Aad Sloomaker, *EMERGO: a generic platform for authoring and playing scenario-based serious games*, OUN

2018-22 Eric Fernandes de Mello Araujo, *Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks*, VUA

2018-23 Kim Schouten, *Semantics-driven Aspect-Based Sentiment Analysis*, EUR

2018-24 Jered Vroon, *Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots*, UT

2018-25 Riste Gligorov, *Serious Games in Audio-Visual Collections*, VUA

2018-26 Roelof Anne Jelle de Vries, *Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology*, UT

2018-27 Maikel Leemans, *Hierarchical Process Mining for Scalable Software Analysis*, TUE

2018-28 Christian Willemsen, *Social Touch Technologies: How they feel and how they make you feel*, UT

2018-29 Yu Gu, *Emotion Recognition from Mandarin Speech*, UVT

2018-30 Wouter Beek, *The "K" in "semantic web" stands for "knowledge": scaling semantics to the web*

—2019—

2019-01 Rob van Eijk, *Web privacy measurement in real-time bidding systems. A graph-based approach to RTB system classification*, UL

2019-02 Emmanuelle Beauxis Aussalet, *Statistics and Visualizations for Assessing Class Size Uncertainty*, CWI, UU

2019-31 Eduardo Gonzalez Lopez de Murillas, *Process Mining on Databases: Extracting Event Data from Real Life Data Sources*, TUE

2019-04 Ridho Rahmadi, *Finding stable causal structures from clinical data*, RUN

2019-05 Sebastiaan van Zelst, *Process Mining with Streaming Data*, TUE

2019-06 Chris Dijkshoorn, *Nichesourcing for Improving Access to Linked Cultural Heritage Datasets*, VU

2019-07 Soude Fazel, *Recommender Systems in Social Learning Platforms*, TUD

2019-08 Frits de Nijs, *Resource-constrained Multi-agent Markov Decision Processes*, TUD

2019-09 Fahimeh Alizadeh Moghaddam, *Self-adaptation for energy efficiency in software systems*, UVA

2019-10 Qing Chuan Ye, *Multi-objective Optimization Methods for Allocation and Prediction*, EUR

2019-11 Yue Zhao, *Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs*, TUD

2019-12 Jacqueline Heinerman, *Better Together*, VU

2019-13 Guanliang Chen, *MOOC Analytics: Learner Modeling and Content Generation*, TUD

2019-14 Daniel Davis, *Large-Scale Learning Analytics: Modeling Learner Behavior & Improving Learning Outcomes in Massive Open Online Courses*, TUD

2019-15 Erwin Walraven, *Planning under Uncertainty in Constrained and Partially Observable Environments*, TUD

2019-16 Guangming Li, *Process Mining based on Object-Centric Behavioral Constraint (OCBC) Models*, TUE

2019-17 Ali Hurriyetoglu, *Extracting actionable information from microtexts*, RUN

2019-18 Gerard Wagenaar, *Artefacts in Agile Team Communication*, UU

- 2019-19 Vincent Koeman, *Tools for Developing Cognitive Agents*, TUD
- 2019-20 Chide Groenouwe, *Fostering technically augmented human collective intelligence*, UU
- 2019-21 Cong Liu, *Software Data Analytics: Architectural Model Discovery and Design Pattern Detection*, TUE
- 2019-22 Martin van den Berg, *Improving IT Decisions with Enterprise Architecture*, VU
- 2019-23 Qin Liu, *Intelligent Control Systems: Learning, Interpreting, Verification*, TUD
- 2019-24 Anca Dumitrache, *Truth in Disagreement - Crowdsourcing Labeled Data for Natural Language Processing*, VU
- 2019-25 Emiel van Miltenburg, *Pragmatic factors in (automatic) image description*, VU
- 2019-26 Prince Singh, *An Integration Platform for Synchronodal Transport*, UT
- 2019-27 Alessandra Antonaci, *The Gamification Design Process applied to (Massive) Open Online Courses*, OUN
- 2019-28 Esther Kuindersma, *Cleared for take-off: Game-based learning to prepare airline pilots for critical situations*, UL
- 2019-29 Daniel Formolo, *Using virtual agents for simulation and training of social skills in safety-critical circumstances*, VU
- 2019-30 Vahid Yazdanpanah, *Multiagent Industrial Symbiosis Systems*, UT
- 2019-31 Milan Jelisavcic, *Alive and Kicking: Baby Steps in Robotics*, VU
- 2019-32 Chiara Sironi, *Monte-Carlo Tree Search for Artificial General Intelligence in Games*, UM
- 2019-33 Anil Yaman, *Evolution of Biologically Inspired Learning in Artificial Neural Networks*, TUE
- 2019-34 Negar Ahmadi, *EEG Microstate and Functional Brain Network Features for Classification of Epilepsy and PNES*, TUE
- 2019-35 Lisa Facey-Shaw, *Gamification with digital badges in learning programming*, OUN
- 2019-36 Kevin Ackermans, *Designing Video-Enhanced Rubrics to Master Complex Skills*, OUN
- 2019-37 Jian Fang, *Database Acceleration on FPGAs*, TUD
- 2019-38 Akos Kadar, *Learning visually grounded and multilingual representations*, OUN
- 2020—
- 2020-01 Armon Toubman, *Calculated Moves: Generating Air Combat Behaviour*, UL
- 2020-02 Marcos de Paula Bueno, *Unraveling Temporal Processes using Probabilistic Graphical Models*, UL
- 2020-03 Mostafa Deghani, *Learning with Imperfect Supervision for Language Understanding*, UvA
- 2020-04 Maarten van Gompel, *Context as Linguistic Bridges*, RUN
- 2020-05 Yulong Pei, *On local and global structure mining*, TUE
- 2020-06 Preethu Rose Anish, *Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support*, UT
- 2020-07 Wim van der Vegt, *Towards a software architecture for reusable game components*, OUN
- 2020-08 Ali Mirsoleimani, *Structured Parallel Programming for Monte Carlo Tree Search*, UL
- 2020-09 Myriam Traub, *Measuring Tool Bias and Improving Data Quality for Digital Humanities Research*, UU
- 2020-10 Alifah Syamsiyah, *In-database Preprocessing for Process Mining*, TU/e
- 2020-11 Sepideh Mesbah, *Semantic-Enhanced Training Data Augmentation Methods for Long-Tail Entity Recognition Models*, TUD
- 2020-12 Ward van Breda, *Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment*, VU
- 2020-13 Marco Virgolin, *Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming*, CWI
- 2020-14 Mark Raasveldt, *Integrating Analytics with Relational Databases*, CWI/UL
- 2020-15 Konstantinos Georgiadis, *Smart CAT: Machine Learning for Configurable Assessments in Serious Games*, OUN
- 2020-16 Ilona Wilmont, *Cognitive Aspects of Conceptual Modelling*, RUN
- 2020-17 Daniele Di Mitri, *The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences*, OUN
- 2020-18 Georgios Methenitis, *Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems*, TUD
- 2020-19 Guido van Capelleveen, *Industrial Symbiosis Recommender Systems*, UT
- 2020-20 Albert Hankel, *Embedding Green ICT Maturity in Organisations*, VU
- 2020-21 Karine da Silva Miras de Araujo, *Where is the robot?: Life as it could be*, VU
- 2020-22 Maryam Masoud Khamis, *Understanding complex systems implementation through a modeling approach: the case of e-government in Zanzibar*, RUN
- 2020-23 Rianne Conijn, *The Keys to Writing: A writing analytics approach to studying writing processes using keystroke logging*, UT
- 2020-24 Lenin da Nobrega Medeiros, *How are you feeling, human? Towards emotionally supportive chatbots*, VUA/RUN
- 2020-25 Xin Du, *The Uncertainty in Exceptional Model Mining*, TUE
- 2020-26 Krzysztof Leszek Sadowski, *GAMBIT: Genetic Algorithm for Model-Based mixed-Integer Optimization*, UU
- 2020-27 Ekaterina Muravyeva, *Personal data and informed consent in an educational context*, TUD
- 2020-28 Bibeg Limbu, *Multimodal interaction for deliberate practice: Training complex skills with augmented reality*, TUD
- 2020-29 Ioan Gabriel Bucur, *Being Bayesian about Causal Inference*, RUN
- 2020-30 Bob Zadok Blok, *Creatief, Creatieve, Creatiefst* UL
- 2020-31 Gongjin Lan, *Learning better - From Baby to Better*, VU
- 2020-32 Jason Rhuggenaath, *Revenue management in online markets: pricing and online advertising*, TUE
- 2020-33 Rick Gilsing, *Supporting service-dominant business model evaluation in the context of business model innovation*, TUE
- 2020-34 Anna Bon, *Intervention or Collaboration? Redesigning Information and Communication Technologies for Development*, MU
- 2020-35 Siamak Farshidi, *Multi-Criteria Decision-Making in*

Software Production, UU

—2021—

- 2021-01** Francisco Xavier Dos Santos Fonseca, *Location-based Games for Social Interaction in Public Space*, TUD
- 2021-02** Rijk Mercur, *Simulating Human Routines: Integrating Social Practice Theory in Agent-Based Models*, TUD
- 2021-03** Seyyed Hadi Hashemi, *Modeling Users Interacting with Smart Devices*, UVA
- 2021-04** Ioana Jivet, *The Dashboard That Loved Me: Designing adaptive learning analytics for self-regulated learning*, OU
- 2021-05** Davide Dell'Anna, *Data-Driven Supervision of Autonomous Systems*, UU
- 2021-06** Daniel Davison, *"Hey robot, what do you think?" How children learn with a social robot*, UT
- 2021-07** Armel Lefebvre, *Research data management for open science*, UU
- 2021-08** Nardie Fanchamps, *The Influence of Sense-Reason-Act Programming on Computational Thinking*, OU
- 2021-09** Cristina Zaga, *The Design of Robothings. Non-Anthropomorphic and Non-Verbal Robots to Promote Children's Collaboration Through Play*, UT
- 2021-10** Quinten Meertens, *Misclassification Bias in Statistical Learning*, UvA
- 2021-11** Anne van Rossum, *Nonparametric Bayesian Methods in Robotic Vision*, UL
- 2021-12** Lei Pi, *External Knowledge Absorption in Chinese SMEs*, UL
- 2021-13** Bob R. Schadenberg, *Robots for Autistic Children: Understanding and Facilitating Predictability for Engagement in Learning*, UT
- 2021-14** Negin Samaeemofrad, *Business Incubators: The Impact of Their Support*, UL
- 2021-15** Onat Ege Adali, *Transformation of Value Propositions into Resource Re-Configurations through the Business Services Paradigm*, TU/e
- 2021-16** Esam A. H. Ghaleb, *BIMODAL EMOTION RECOGNITION FROM AUDIO-VISUAL CUES*, UM
- 2021-17** Dario Dotti, *Human Behavior Understanding from motion and bodily cues using deep neural networks*, UM
- 2021-8** Remi Wieten, *Bridging the Gap Between Informal Sense-Making Tools and Formal Systems - Facilitating the Construction of Bayesian Networks and Argumentation Frameworks*, UU
- 2021-19** Roberto Verdecchia, *Architectural Technical Debt: Identification and Management*, VU
- 2021-20** Masoud Mansoury, *Understanding and Mitigating Multi-Sided Exposure Bias in Recommender Systems*, TU/e
- 2021-21** Pedro Thiago Timbó Holanda, *Progressive Indexes*, CWI
- 2021-22** Sihang Qiu, *Conversational Crowdsourcing*, TUD
- 2021-23** Hugo Manuel Proença, *Robust rules for prediction and description*, LIACS
- 2021-24** Kaijie Zhu, *On Efficient Temporal Subgraph Query Processing*, TUE
- 2021-25** Eoin Martino Grua, *The Future of E-Health is Mobile: Combining AI and Self-Adaptation to Create Adaptive E-Health Mobile Applications*, VUA
- 2021-26** Benno Kruit, *Reading the Grid: Extending Knowledge Bases from Human-readable Tables*, CWI & VUA
- 2021-27** Jelte van Waterschoot, *Personalized and Personal Conversations: Designing Agents Who Want to Connect With You*, UT
- 2021-28** Christoph Selig, *Understanding the Heterogeneity of Corporate Entrepreneurship Programs*, UL

—2022—

- 2022-1** Judith van Stegeren), *Flavor text generation for role-playing video games*, UT
- 2022-2** Paulo da Costa, *Data-driven Prognostics and Logistics Optimisation: A Deep Learning Journey*, TU/e
- 2022-3** Ali el Hassouni, *A Model A Day Keeps The Doctor Away: Reinforcement Learning For Personalized Healthcare*, VUA
- 2022-4** Ünal Aksu, *A Cross-Organizational Process Mining Framework*, UU
- 2022-5** Shiwei Liu, *Sparse Neural Network Training with In-Time Over-Parameterization*, TU/e
- 2022-6** Reza Refaei Afshar, *Machine Learning for Ad Publishers in Real Time Bidding*, TU/e
- 2022-7** Sambit Prahara, *Measuring the Unmeasurable? Towards Automatic Co-located Collaboration Analytics*, OU
- 2022-8** Maikel L. van Eck, *Process Mining for Smart Product Design*, TU/e
- 2022-9** Oana Andreea Inel, *Understanding Events: A Diversity-driven Human-Machine Approach*, VUA
- 2022-10** Felipe Moraes Gomes, *Examining the Effectiveness of Collaborative Search Engines*, TUD
- 2022-11** Mirjam de Haas, *Staying engaged in child-robot interaction, a quantitative approach to studying preschoolers' engagement with robots and tasks during second-language tutoring*, UT
- 2022-12** Guanyi Chen, *Computational Generation of Chinese Noun Phrases*, UU
- 2022-13** Xander Wilcke, *Machine Learning on Multimodal Knowledge Graphs: Opportunities, Challenges, and Methods for Learning on Real-World Heterogeneous and Spatially-Oriented*, VUA
- 2022-14** Michiel Overeem, *Evolution of Low-Code Platforms*, UU
- 2022-15** Jelmer Jan Koorn, *Work in Process: Unearthing Meaning using Process Mining*, UU
- 2022-16** Pieter Gijbbers, *Systems for AutoML Research*, TU/e
- 2022-17** Laura van der Lubbe, *Empowering vulnerable people with serious games and gamification*, VUA
- 2022-18** Paris Mavromoustakos Blom, *Player Affect Modelling and Video Game Personalisation*, TU
- 2022-19** Bilge Yigit Ozkan, *Cybersecurity Maturity Assessment and Standardisation*, UU
- 2022-20** Fakhra Jabeen, *Dark Side of the Digital Media - Computational Analysis of Negative Human Behaviors on Social Media*, VUA
- 2022-21** Seethu Mariyam Christopher, *Intelligent Toys for Physical and Cognitive Assessments*, UM
- 2022-22** Alexandra Sierra Rativa, *Virtual Character Design and its potential to foster Empathy, Immersion, and Collaboration Skills in Video Games and Virtual Reality Simulations*, TIU
- 2022-23** Ilir Kola, *Enabling Social Situation Awareness in Support Agents*, TUD
- 2022-24** Samaneh Heidari, *Agents with Social Norms and Values - A framework for agent based social simulations with social norms and personal values*, UU
- 2022-25** Anna L.D. Latour, *Optimal decision-making under constraints and uncertainty*, LU
- 2022-26** Anne Dirkson, *Knowledge Discovery from Patient Forums: Gaining novel medical insights from patient experiences*, LU
- 2022-27** Christos Athanasiadis, *Emotion-aware cross-modal domain adaptation in video sequences*, UM

2022-28 Onuralp Ulusoy, *Privacy in Collaborative Systems*, UU

2022-29 Jan Kolkmeier, *From Head Transform to Mind Transplant: Social Interactions in Mixed Reality*, UT

2022-30 Dean De Leo, *Analysis of Dynamic Graphs on Sparse Arrays*, CWI

2022-31 Konstantinos Traganos, *Tackling Complexity in Smart Manufacturing with Advanced Manufacturing Process Management*, TU/e

2022-32 Cezara Pastrav, *Social simulation for socio-ecological systems*, UU

2022-33 Brinn Hekkelman, *Fair Mechanisms for Smart Grid Congestion Management*, TUD/CWI

2022-34 Nimat Ullah, *Mind Your Behaviour: Computational Modelling of Emotion & Desire Regulation for Behaviour Change*, VUA

2022-35 Mike E.U. Ligthart, *Shaping the Child-Robot Relationship: Interaction Design Patterns for a Sustainable Interaction*, VUA

—2023—

2023-1 Bojan Simoski, *Untangling the Puzzle of Digital Health Interventions*, VUA

2023-2 Mariana Rachel Dias da Silva, *Grounded or in flight? What our bodies can tell us about the whereabouts of our thoughts*, TiU