

An Argumentative Framework for Generating Explainable Group Recommendations

Citation for published version (APA):

van Kuijk, K., Mahmoudi, S. S., Wen, Y., Barile, F., & Rienstra, T. (2023). An Argumentative Framework for Generating Explainable Group Recommendations. In *Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization* (pp. 266–274). Association for Computing Machinery. <https://doi.org/10.1145/3563359.3597387>

Document status and date:

Published: 01/01/2023

DOI:

[10.1145/3563359.3597387](https://doi.org/10.1145/3563359.3597387)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



An Argumentative Framework for Generating Explainable Group Recommendations

Kristian van Kuijk

Maastricht University

Maastricht, Netherlands

kristian.vankuijk@student.maastrichtuniversity.nl

Yangfan Wen

Maastricht University

Maastricht, Netherlands

yangfan.wen@student.maastrichtuniversity.nl

Seyedeh Sara Mahmoudi

Maastricht University

Maastricht, Netherlands

s.mahmoudi@student.maastrichtuniversity.nl

Francesco Barile

Maastricht University

Maastricht, Netherlands

f.barile@maastrichtuniversity.nl

Tjitze Rienstra

Maastricht University

Maastricht, Netherlands

t.rienstra@maastrichtuniversity.nl

ABSTRACT

In the context of group recommender systems, explanations strategies have been proposed to improve recommendations perceived fairness, consensus, satisfaction, and to help the group members in the decision-making process. In general, such explanations try to clarify the underlying social choice-based aggregation strategies used to generate the recommendations. However, results in the literature are conflicting, and the real benefit of such explanations seem to be limited. In this work, we propose a novel approach, which makes use of an argumentative framework built using information about the aspects that are connected to the recommended items. Such framework is used to generate recommendations, and related explanations. We provide a proof of concept on how to generate explanations for the group, as well as specific explanations for the group members, which use the information in the argumentative frameworks to enrich the explanations. Furthermore, we propose privacy-preserving versions for the explanations, as well as a graphical approach based on tag clouds. In future works, we plan to evaluate the quality of the provided recommendations in offline settings, as well as the impact of the proposed explanations in a series of user studies.

CCS CONCEPTS

• Information systems → Recommender systems; Decision support systems; Social recommendation; • Computing methodologies → Knowledge representation and reasoning.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP '23 Adjunct, June 26–29, 2023, Limassol, Cyprus

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9891-6/23/06...\$15.00

<https://doi.org/10.1145/3563359.3597387>

KEYWORDS

Explainable AI, Explainable Recommender Systems, Group Recommender Systems, Argumentative Framework

ACM Reference Format:

Kristian van Kuijk, Seyedeh Sara Mahmoudi, Yangfan Wen, Francesco Barile, and Tjitze Rienstra. 2023. An Argumentative Framework for Generating Explainable Group Recommendations. In *Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization (UMAP '23 Adjunct)*, June 26–29, 2023, Limassol, Cyprus. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3563359.3597387>

1 INTRODUCTION

Recommender Systems (RS) are purpose-built to provide recommendations to users by analyzing their interests, preferences, and past behaviors, while also filtering through the vast amounts of data to identify relevant items. A Group Recommender System (GRS) uses the preferences and interactions of a group of users to provide customized recommendations for the entire group, taking into account the unique characteristics of the group members and the context of the recommendation task. Commonly, this task is performed by aggregating the individual group members' preferences or recommendations [36], using *social choice-based aggregation strategies* [23, 24].

In the context of Recommender Systems, explanations have been used to help improve the transparency, persuasiveness, effectiveness, trustworthiness, and satisfaction in RS [12, 47], providing intuitive explanations accompanying users personalized recommended items [21, 38]. When considering group recommendations, explanations may also be used to help users agree on a joint decision, as well as improve users' perceived fairness, perceived consensus, and satisfaction [2, 12, 29, 43]. In general, Explainable Group Recommender Systems (XGRS) attempt to offer comprehensible and interpretable justifications for recommendations given to a groups of users, on the basis of the underlying aggregation strategy used to generate the group recommendation [2, 43].

The use of XGRS can support interactions among users, create group cohesion, and simplify decision-making in groups [12].

However, the state-of-the-art explanations only provide very few information within the explanations [19, 29, 43], and it may be difficult for users to understand the reasoning behind the suggested items and the methods used to generate them.

In the context of GRS, another aspect to consider when providing explanations is the necessity of preserving the user's privacy, as the information included in the explanations are shared with all the group members [26–28]. Privacy in GRS refers to protecting the sensitive information of individual members within the group. This usually includes personal preferences and ratings, even though other information could potentially be used to identify them. Privacy concerns are particularly important in group recommendation systems since they often involve sharing information between multiple users, increasing the risk of sensitive information being exposed or misused.

Recently, argumentative frameworks have been proposed to implement explainable recommender systems capable of providing accurate recommendations with informative explanations, showing promising results [35]. An argumentative framework is a structure modelling a reasoning process in terms of arguments and relationships between them, such as one argument attacking or supporting another one [10]. Given such a model, conclusions are drawn by determining which arguments win the “debate” represented by the argumentation framework.

This paper proposes an argumentative framework for GRS to generate explainable decisions for groups. More specifically, we extend the approach proposed by Rago et al. [35], originally designed for single user RS, to Group Recommender Systems. We apply our approach on three widely used aggregation strategies, a consensus based, Average Satisfaction (AVG), and two borderline strategies, Least Misery (LM) and Most Pleasure (MP) [24, 36] (see Section 2.1). Furthermore, we propose strategies to generate explanations for the group recommendations which can use the information provided by the argumentative framework both on a group and also on an individual level, considering basic and also privacy-preserving strategies.

The remainder of the paper is organized as follows: in the Section 2 we provide an overview of related work in the field of group recommender systems, explainable group recommender systems, and argumentative frameworks, describing different types of techniques and frameworks used in these areas. In Section 2.4, we present the proposed argumentative group recommender system that combines argumentation theory and group decision-making to provide explainable recommendations to groups by generating a set of arguments for and against each recommended item. Also, we propose a strategies to generate explanations, using an argumentative framework to provide privacy-aware explanations to users, and visualizing the explanations using tag-cloud representation. Finally, in the Section 4, we illustrate conclusions and future works.

2 RELATED WORK

In this section, we introduce the basic aggregation strategies used to generate group recommendations. Furthermore, we explore the main works on explainability for group recommender systems, and we illustrate a novel line of research evaluating privacy concerns related to personal information included in group recommendations.

Finally, we introduce argumentative frameworks for recommender systems.

2.1 Social Choice-based Aggregation Strategies

Group Recommender Systems (GRSs) are normally implemented using two main strategies: (i) *models aggregation*, where individual preferences are aggregated to create a group model, which is then used to generate the group recommendations; and (ii) *predictions aggregation*, where individual predictions and recommended items are aggregated, and the items with the highest aggregated scores are suggested to the group [12]. In both cases, the aggregation strategies inspired by *Social Choice Theory* - the study of systems for making collective choices which affect a group of people [20] - are often used to perform such aggregations [24]. In this work, we rely on variations of three widely used *social choice-based aggregation strategies*: *Average (AVG)*, a *consensus-based* strategy that recommends the item with the highest average of all group members' ratings [36]; *Least Misery (LM)* and *Most Pleasure (MP)* which are *borderline* strategies. LM recommends the item which has the highest of all lowest ratings, while MP suggests the item with the highest individual group member rating [36]. Several studies have been performed to compare the performances of the different strategies. Masthoff and Delić [25] show that different strategies perform better than others in two different experimental settings in terms of perceived group satisfaction. Felfernig et al. [11] shows that in real-life scenario, the preferred preference aggregation procedures differ from low-involvement item domains (like restaurants and movies) and high-involvement item domains (such as decisions on new automobiles, financial services, and housing). The latter are frequently the subjects of repetitive group decisions (e.g., the same group selects a restaurant for a dinner every three months). In high involvement item domains, groups frequently use LM methods, while AVG is preferred in low involvement item domains. Although we use a low-investment domain (movies), our proposed approach aims at being used on different domains; hence we decided to present strategies based on both AVG and LM, together with a strategy based on MP.

2.2 Explainable Group Recommender Systems

Explanations has been used in recommender systems to achieve different goals [16, 41]. Earlier studies have shown that the use of explanations led to an increased quality, decision support, trust, overall satisfaction and higher acceptance of the recommendations [33, 40]. Several strategies has been proposed to increase transparency, also combining different explanation styles. However, for many recommendation systems it is still difficult to adequately explain to customers how and why particular suggestions are made [37]. In general, explanations rely on the underlying recommendation strategy, using the information that have been important in determining the systems' suggestion. Content-based recommender systems suggest items to users based on their preferences, by calculating the similarity between item's on the basis of their content information. The possibility to link recommended items with similar items the user consumed in past interactions provides a strong explanatory power to these strategies. One example is provided in Vig et al. [45] where tagsplanations are proposed: tags associated

to contents are used to generate explanations to describe certain movies or recipes. On the contrary, collaborative filtering methods may be difficult to explain to users since the patterns which generated the recommendations may not be easily interpretable. Whether it is the matrix factorization model [22] with high prediction accuracy, the neural network-based collaborative filtering [7] with strong versatility, or the flexible latent factor model [17], all fall short in terms of explainability due to the complexity of the underlying model. To overcome this issue, there has been numerous attempts to improve transparency through the use of visual aids including flowcharts, clustermaps and map visualizations [18, 44], or hierarchical visualization [46]. Collaborative filtering and content-based methods have both been combined in hybrid approaches. These methods, however, frequently still fall short in terms of clearly outlining the suggestion process to users.

In group recommendations, explanations have been used to achieve additional goals, such as improving the fairness within the group, or helping the group members agree on the decision [12]. However, only a few studies investigate the problem of generating explanations for groups. Typically, such explanations are related to the underlying mechanism of the employed social choice-based aggregation strategy [19, 29, 43]. Najafian and Tintarev [29] introduced textual explanations that could be reassuring, in case of an agreement in the group, or repairing in the case of a disagreement. Kapcak et al. [19] extended this work using the wisdom of the crowd to improve the quality of the initially proposed explanations. In Quijano-Sanchez et al. [34], factors like group members' personality and tie strength between them are used to generate tactful explanations, in order to avoid damaging friendships. Tran et al. [43] proposed a user study to evaluate explanations for six social choice-based aggregation strategies in terms of *fairness perception* and *consensus perception*, and user *satisfaction* regarding the group recommendation. Barile et al. [2] proposed a reproduction of this study, but the results show no significant effect from the presence of an explanation, when compared to a control condition without explanations. Nguyen and Ricci [31] introduced a utility vector in GRS, for both group members and the whole group. The utility vector indicates the importance of each aspect for the individual group members and the whole group. However, our approach provides explicit attack and support relationship through the tripolar argumentation framework connecting aspects to different items (see Section 3), allowing to produce more expressive explanations connecting also different items related to the same aspects. In this work, we present an approach aiming at providing recommendations for groups with comparable accuracy with respect to the baseline strategies in the state of the art, but with an higher explanatory power.

2.3 Privacy in Group Recommender Systems Explanations

When providing explanations for groups, another goal that should be considered is the privacy-preservation [26–28]. Privacy protection is a challenging issue, as a simple strategy to improve the explanation of recommended results is to inform group members about other people's preferences. Nevertheless, some users might not accept to share information that could be really personal with

the other group members. [19] highlighted how some group members appreciated giving more details about ratings in the explanations, while some on the other hand had concerns about privacy violation. [26] investigates privacy concerns regarding group recommendation explanation in respect to personality, relationship type or preference scenario. They evaluated the impact on privacy concerns related to several types of information in a user study. They showed significant effects related to the specific group preference scenario (whether the participant was in minority or a majority position in the group) and type of relationship between the group members. Furthermore, they showed an impact of personality, as some personality factors (Agreeableness and Extraversion, from the Five Factors Model [15]) were positively correlated with privacy concerns. In this work, we present different strategies for generating explanations for group recommendations, proposing for each a base definition, and also a privacy-preserving alternative.

2.4 Argumentative Frameworks for Recommender Systems

A number of abstract argumentation frameworks exist [10], ranging from bipolar argumentation frameworks [8], tripolar frameworks [14], and generalized frameworks [3]. The RS proposed by [35], which this paper is based on, can be considered as a special instance of the generalized framework. Previous literature has utilized defeasible logic programming to enhance argument-based analysis techniques [6, 9, 39], repaired recommendations using rule-based arguments in user interactions in hybrid RS [4], and simulated explanation-based argument models based on supporting relations [30, 42]. Compared to the aforementioned literature, the most significant advantage of [35] RS is its interpretability, while maintaining effectiveness (precision) without sacrificing interpretability. Furthermore, this RS is also scalable and supports feedback mechanisms, and its argumentation framework allows for the extraction of various explanations. Hence, we propose an extension of the approach proposed in [35] which can be used to generate recommendations for a group of users, and use the information from the underlying tripolar framework to generate explanations for the group and for the individual group members, together with alternative privacy-preserving formulations.

3 AN ARGUMENTATIVE FRAMEWORK FOR GENERATING EXPLAINABLE GROUP RECOMMENDATIONS

In this section we present our argumentative approach for generating group recommendations and explanations. The approach is based on the argumentative approach to generating explainable recommendations due to Rago et al. [35], which we outline in section 3.1. We then explain how we extend this approach to generate group recommendations. Finally, we discuss several strategies to generate explanations, for the group as a whole and for the single group members, also introducing privacy preserving versions of our explanations, together with a graphical approach (based on tag clouds).

3.1 An Argumentative Recommender System

Our formal model extends the argumentation-based approach to generating explainable recommendations due to Rago et al. [35]. This approach is based on two components. The first is an *aspect-item framework*, which is used to represent known ratings of movies and their aspects by users, and to predict unknown ratings. The second is a *tripolar argumentation framework*, which is generated on the basis of an aspect-item framework, and is used to generate argumentative explanations for recommendations. In this section we formally define these two notions, and then explain how we extend the approach to provide group recommendations.

An aspect-item framework consists of a set $\mathcal{I}$ of items (e.g., movies) to be recommended, a set $\mathcal{T}$ of aspect types (e.g., genre, director, actor) and a set $\mathcal{A}$ of aspects, each belonging to a type (e.g., the aspect “comedy” of type genre, the aspect “Robert Zemeckis” is of type director, and so on). Items are linked with their aspects by the relation $\mathcal{L}$. For each user, the partial function $\mathcal{R}$ associates items and aspects with known ratings, which are real numbers in the $[-1, 1]$ interval.

Definition 3.1. An *aspect-item (A-I) framework* is a tuple $\langle \mathcal{I}, \mathcal{A}, \mathcal{T}, \mathcal{L}, \mathcal{U}, \mathcal{R} \rangle$ where:

- Disjoint, finite and non-empty sets $\mathcal{I}$ and $\mathcal{A}$ containing, respectively, items and aspects. We use X to denote $\mathcal{I} \cup \mathcal{A}$
- $\mathcal{T}$ a finite non-empty set of types, such that for each aspect $a \in \mathcal{A}$, there is a unique $t \in \mathcal{T}$ that is the type of a . We use $\mathcal{A}_t$ to denote the set aspects of type t .
- $\mathcal{L} \subseteq (\mathcal{I} \times \mathcal{A}) \cup (\mathcal{A} \times \mathcal{I})$ a symmetric binary relation. We use $\mathcal{L}(x)$ to denote $\{y \in X \mid (y, x) \in \mathcal{L}\}$ and $\mathcal{L}_t(i)$ to denote $\{a \in \mathcal{L}(i) \mid a \in \mathcal{A}_t\}$;
- $\mathcal{U}$ a finite non empty-set of users;
- $\mathcal{R} : \mathcal{U} \times X \rightarrow [-1, 1]$ a partial function of known ratings. If the rating of user u for item i is unknown then $\mathcal{R}(u, i)$ is undefined.

Example 3.2. As an example, consider the A-I framework $\langle \mathcal{I}, \mathcal{A}, \mathcal{T}, \mathcal{L}, \mathcal{U}, \mathcal{R} \rangle$ where:

- $\mathcal{I} = \{i_1, i_2\}$.
- $\mathcal{A} = \{a_1, d_1, d_2, g_1, g_2\}$.
- $\mathcal{T} = \{\text{actor}, \text{director}, \text{genre}\}$ with $\mathcal{A}_{\text{actor}} = \{a_1\}$, $\mathcal{A}_{\text{director}} = \{d_1, d_2\}$ and $\mathcal{A}_{\text{genre}} = \{g_1, g_2\}$.
- $\mathcal{L} = \{(i_1, a_1), (i_1, d_1), (i_1, g_1), (i_1, g_2), (i_2, a_1), (i_2, d_2), (i_2, g_2)\}$.
- $\mathcal{U} = \{u_1, \dots, u_n\}$.
- $\mathcal{R}$ defines the known ratings $\mathcal{R}(u_1)(a_1) = -0.6$, $\mathcal{R}(u_1)(d_1) = 0.4$, and $\mathcal{R}(u_1)(g_2) = 0.5$. The rating $\mathcal{R}(u_1)(i_1)$ is unknown and thus undefined.

Figure 1 includes a graphical representation of this A-I framework.

We predict unknown ratings using the same approach taken by Rago et al. [35]. Our approach is somewhat simpler since we do not consider variable degrees of importance for collaborative filtering or aspect types. We determine the predicted rating using the function $\mathcal{P}_{\mathcal{I}}^u : \mathcal{I} \rightarrow [-1, 1]$ defined as follows.

$$\mathcal{P}_{\mathcal{I}}^u(i) = \begin{cases} \mathcal{R}(u, i) & \text{if } \mathcal{R}(u, i) \text{ is defined,} \\ \frac{\sum_{t \in \mathcal{T}} [\sum_{a \in \mathcal{L}_t(i)} \mathcal{P}_{\mathcal{A}}^u(a)] / |\mathcal{L}_t(i)|}{|\mathcal{T}|} & \text{otherwise.} \end{cases}$$

where $\mathcal{P}_{\mathcal{A}}^u$ is the predicted aspect rating defined below. Put simply, the predicted rating is the average rating on the item from similar

users and the aspect ratings from each of the linked aspects (in the case the rating of the user is unknown). Note that Rago et al. [35] originally define a profile π^u per user $u \in \mathcal{U}$ and cases in which the aspect ratings might be undefined. For simplicity, we skip those steps in this paper. Similarly to items, the predicted aspect rating is determined by the function $\mathcal{P}_{\mathcal{A}}^u : \mathcal{A} \rightarrow [-1, 1]$ defined as follows.

$$\mathcal{P}_{\mathcal{A}}^u(a) = \begin{cases} \mathcal{R}(u, a) & \text{if } \mathcal{R}(u, a) \text{ is defined,} \\ \left(\frac{\sum_{i \in \Lambda^u(a)} \mathcal{R}(u, i)}{|\Lambda^u(a)|} + \frac{\sum_{i \in \Lambda^{-u}(a)} \mathcal{P}^u(i)}{|\Lambda^{-u}(a)|} \right) / 2 & \text{otherwise.} \end{cases}$$

where $\Lambda^u(a)$ the set of linked items with defined ratings from u , $\Lambda^{-u}(a)$ the set of linked items without any ratings from u , and $\mathcal{P}^u(i)$ the rating from similar users to u for i .

Example 3.3. (Continued from Example 3.2) In Example 3.2, the rating of user u_1 on item i_1 is unknown. The prediction of this rating is the average of the known ratings of the linked aspects:

$$\mathcal{P}_{\mathcal{I}}^{u_1}(i_1) = \frac{\mathcal{P}_{\mathcal{A}}^{u_1}(a_1) + \mathcal{P}_{\mathcal{A}}^{u_1}(d_1) + \mathcal{P}_{\mathcal{A}}^{u_1}(g_1)}{|\mathcal{T}|} = \frac{-0.6 + 0.4 + 0.5}{3} = 0.1$$

In this example, the ratings of all aspects linked with i_1 are known. If, for instance, the rating of user u_1 on a_1 is not known, but the rating of user u_1 on i_2 is, then the rating of a_1 would be computed based on the rating of i_2 . If the rating of user u_1 for these items is not known, but the ratings of other users are, then the rating can be determined via the rating $\mathcal{P}^{u_1}(i_2)$ from similar users to u_1 on i_2 .

A *tripolar argumentation (TF) framework* consists of a set X of arguments and three dialectical relationships among arguments: $\mathcal{L}^-$ (attack), $\mathcal{L}^+$ (support), and $\mathcal{L}^0$ (neutralisation). We generate a TF on the basis of an A-I framework and a user. This TF will be used to extract explanations for recommendations. We first define the notion of TF formally and then define how a TF is generated.

Definition 3.4. A *tripolar argumentation framework (TF)* is a tuple $\langle X, \mathcal{L}^-, \mathcal{L}^+, \mathcal{L}^0 \rangle$ where X is a set of arguments and $\mathcal{L}^-, \mathcal{L}^+, \mathcal{L}^0$ are binary relations over X . For $x, y \in X$, we say that x attacks y if $(x, y) \in \mathcal{L}^-$, x supports y if $(x, y) \in \mathcal{L}^+$, and x neutralises y if $(x, y) \in \mathcal{L}^0$. We use $L^\times(x)$ to denote $\{y \in X \mid (y, x) \in \mathcal{L}^\times\}$ the attackers, supporters or neutralisers of x .

Figure 1 shows a graph-based representation of an aspect-item framework, with nodes representing arguments and items, and edges representing dialectical relationships. The first step in generating a TF for a given user and A-I framework, is to assign a direction to the links of the A-I framework, such that an item/aspect x points to another item/aspect y whenever the rating of x affects the rating of y . For instance, if a rating by user u of an item/aspect is known then this item/aspect has no incoming links. We refer to the result of assigning a direction to the links as a *directed A-I*.

Definition 3.5. Given an A-I $\langle \mathcal{I}, \mathcal{A}, \mathcal{T}, \mathcal{L}, \mathcal{U}, \mathcal{R} \rangle$ and user $u \in \mathcal{U}$, the *directed A-I* for u is the A-I $\mathcal{F}^u = \langle \mathcal{I}, \mathcal{A}, \mathcal{T}, \mathcal{L}^u, \mathcal{U}, \mathcal{R} \rangle$ where $\mathcal{L}^u = \{(i, a) \in \mathcal{L} \mid \mathcal{R}(u, a) \text{ is undefined and } \exists v \in \mathcal{U} \text{ s.t. } \mathcal{R}(v, i) \text{ is defined}\} \cup \{(a, i) \in \mathcal{L} \mid \mathcal{R}(u, i) \text{ is undefined}\}$. For $x \in X$, we refer to $\mathcal{L}^u(x) = \{y \in X \mid (y, x) \in \mathcal{L}^u\}$ as the set of item-aspects *affecting* x . Also, for $i \in \mathcal{I}$, we use $\mathcal{L}_t^u(i)$ to denote the set $\{a \in \mathcal{L}^u(i) \mid a \in \mathcal{A}_t\}$.

The next step is to convert the directed A-I $\mathcal{F}^u$ into a TF. This TF consists of all items and aspects, and consists of relationships (support, attack, neutralisation) based on the polarity of the links in $\mathcal{L}_t^u$:

Definition 3.6. The TF corresponding to the directed A-I $\mathcal{F}^u = \langle \mathcal{I}, \mathcal{A}, \mathcal{T}, \mathcal{L}^u, \mathcal{U}, \mathcal{R} \rangle$ is the TF $\langle \mathcal{X}, \mathcal{L}^-, \mathcal{L}^+, \mathcal{L}^0 \rangle$ where:

- $\mathcal{L}^- = \{(i, a) \in \mathcal{L}^u | r^u(i) < 0\} \cup \{(a, i) \in \mathcal{L}^u | \mathcal{P}_{\mathcal{A}}^u(a) < 0\}$;
- $\mathcal{L}^+ = \{(i, a) \in \mathcal{L}^u | r^u(i) > 0\} \cup \{(a, i) \in \mathcal{L}^u | \mathcal{P}_{\mathcal{A}}^u(a) > 0\}$;
- $\mathcal{L}^0 = \{(i, a) \in \mathcal{L}^u | r^u(i) = 0\} \cup \{(a, i) \in \mathcal{L}^u | \mathcal{P}_{\mathcal{A}}^u(a) = 0\}$;

The TF corresponding the direct A-I $\mathcal{F}^u$ explains the items recommended to user u by stating which aspects and items positively or negatively affect other aspects and items. One can thus find reasons for whether a given item is or is not recommended. In the next section we will explain how we extend this method to group recommendations. In Section 3.3 we will explain how we extract actual explanations from a TF.

Example 3.7. (Continued from Example 3.3) From the ratings shown in Example 3.2, we get the TF $\langle \mathcal{X}, \mathcal{L}^-, \mathcal{L}^+, \mathcal{L}^0 \rangle$ where $\mathcal{L}^- = \{(a_1, i_1)\}$, $\mathcal{L}^+ = \{(d_1, i_1), (g_1, i_1)\}$ and $\mathcal{L}^0 = \emptyset$. This TF demonstrates that aspect a_1 has a negative influence on the rating of i_1 , while d_1 and g_1 have a positive influence.

3.2 An Argumentative Group Recommender System

Our approach to extend the argumentative recommender system for providing group recommendations can be seen as a *predictions aggregation* approach [12], as the general idea is to aggregate the predicted aspect ratings $\mathcal{P}_{\mathcal{A}}^u(a)$ for the group members to obtain the the predicted group ratings, and then recommend the items with the highest aggregated scores. However, in order to use the TF as defined in the Definition 3.6, we introduce virtual users representing each group. More formally, we define a set of groups $g_1, \dots, g_n$, where $g_1 \cup g_2 \cup \dots \cup g_n \in \mathcal{U}$. In principle, a generic user can belong to more groups, or to no groups. We define three approaches based on three widely-used social-choice aggregation strategies: (i) average aggregation (AVG); (ii) least misery (LM); and (iii) most pleasure (MP).

Average Aggregation (AVG). For each group $g = u_1, u_2, \dots, u_k$, we introduce a virtual user u_g representing the group, and we compute:

- $\mathcal{P}_{\mathcal{I}}^{u_g}(i) = \sum_{u \in g} \mathcal{P}_{\mathcal{I}}^u(i) / |g|$
- $\mathcal{P}_{\mathcal{A}}^{u_g}(a) = \sum_{u \in g} \mathcal{P}_{\mathcal{A}}^u(a) / |g|$

Least Misery (LM). For each group $g = u_1, u_2, \dots, u_k$, we introduce a virtual user u_g representing the group, and we compute:

- $\mathcal{P}_{\mathcal{I}}^{u_g}(i) = \min_{u \in g} \mathcal{P}_{\mathcal{I}}^u(i)$
- $\mathcal{P}_{\mathcal{A}}^{u_g}(a) = \min_{u \in g} \mathcal{P}_{\mathcal{A}}^u(a)$

Most Pleasure (MP). For each group $g = \{u_1, u_2, \dots, u_k\}$, we introduce a virtual user u_g representing the group, and we compute:

- $\mathcal{P}_{\mathcal{I}}^{u_g}(i) = \max_{u \in g} \mathcal{P}_{\mathcal{I}}^u(i)$
- $\mathcal{P}_{\mathcal{A}}^{u_g}(a) = \max_{u \in g} \mathcal{P}_{\mathcal{A}}^u(a)$

Group tripolar argumentation framework. Once we chose an aggregation strategy, we can generate a TF for each group g using the A-I of the corresponding virtual user u_g and the definition 3.6. The group TF will be used for generating group explanations. We illustrate a complete overview of the pipeline in Fig. 1.

3.3 Generating Explanations for Group Recommendations

We propose a different approach to Rago et al. [35] in providing explanations to recommendations. Felfernig et al. [13] separate explanations in two categories of recommender systems: collaborative filtering and content based. We use a similar approach: a first part of the explanation refers to the aggregation strategy, while we use the aspects from the A-I framework as the content based part of the system.

3.3.1 Explaining Social Choice-based Aggregations. Our explanations generated by taking the aggregation functions into account are:

- **Least Misery:** “item y has a group score of x due to the (lowest) rating determined for user a ”.
- **Most Pleasure:** “item y has a group score of x due to the (highest) rating determined for user b ”.
- **Average:** “item y is most similar to the ratings of users a , b , and c ”.

We can notice that none of those explanations take privacy into account. For privacy reason, in a group-setting the rating of individual users should not be available in the explanations. Taking privacy into account, we obtain the following privacy-preserving explanations:

- **Least Misery:** “item y is recommended because it avoids misery within the group”.
- **Most Pleasure:** “item y is recommended because it leads to most pleasure within the group”.
- **Average:** “item y is most similar to the ratings of users of the group”.

3.3.2 Content-based Explanations. We also defined texts to be used to enrich the basic explanations and provide information about the content of the recommender items. Following Felfernig et al. [13], we formulate the following explanation: “Item t is recommended since each group member is interested in aspect a ”. Similarly as before, taking privacy into account, we generate the explanation: “Item t is recommended since each group member is interested in aspect a ”. These general examples can, in principles, be implemented for any aspects related to the recommended items. However, previous studies Najafian and Tintarev [29] highlighted how users prefer short, simple, informal and friendly explanations. Hence, it is important to only select the most important aspects to be considered in the explanation. Our proposed strategy is to prioritize extreme ratings (i.e. worst and best rated aspects). We typically choose one or two aspects that are supporting of the selected item in the group TF, selecting the one or the two best rated aspects from the A-I framework. Examples of group explanations for each considered aggregation strategy are illustrated in the Table 1.

3.3.3 Individual Explanations for Group Members. While so far we only focused on generating explanations to the group as a whole, using the information obtained by group TF, explanations can also be delivered to the individual users, using their individual TF. In this case, privacy concerns of the specific user should not be considered




Figure 1: An overview of the GRS pipeline with both the aspect-item framework and the tripolar argumentation framework. Similarly to Rago et al. [35], we defined each node label in the aspect-item framework as $(x\text{'s name}, \mathcal{R}(u, x), \mathcal{R}(v, x), \mathcal{P}_X^u(x))$, with $\mathcal{U} = \{u, v\}$. In the figure, squares refer to movies, circles to actors, triangles to movie genres and parallelograms to directors. In the tripolar argumentation framework, '+' denotes support, '-' attack and '0' neutral.

Strategy	Linguistic Explanation
LM	Catch me if you can is recommended because it avoids misery within the group. Furthermore, the movie is recommended since the group as a whole is interested in <i>Drama</i> and <i>Tom Hanks</i> .
MP	Catch me if you can is recommended because it leads to most pleasure within the group. Furthermore, the movie is recommended since the group as a whole is interested in <i>Drama</i> and <i>Tom Hanks</i> .
AVG	Catch me if you can is most similar to the ratings of users of the group. Furthermore, the movie is recommended since the group as a whole is interested in <i>Drama</i> and <i>Tom Hanks</i> .

Table 1: Group explanation examples for LM, MP and AVG given the recommended item *Catch me if you can* generated from the group TF. Items are denoted in bold and aspects in italic.

Strategy	Linguistic Explanation
LM	Catch me if you can is most similar to the ratings of users of the group, due to the lowest rating determined for you. Furthermore, the movie is recommended since the group as a whole is interested in <i>Drama</i> and <i>Tom Hanks</i> , and since you particularly like movies featuring <i>Leonardo di Caprio</i> .
MP	Catch me if you can is most similar to the ratings of users of the group, with the highest rating determined for you. Furthermore, the movie is recommended since the group as a whole is interested in <i>Drama</i> and <i>Tom Hanks</i> , and since you particularly like movies featuring <i>Leonardo di Caprio</i> .
AVG	Catch me if you can is most similar to the ratings of users of the group, despite a predicted rating of 2.9 for you. Furthermore, the movie is recommended since the group as a whole is interested in <i>Drama</i> and <i>Tom Hanks</i> , and since you particularly like movies featuring <i>Leonardo di Caprio</i> .

Table 2: Group explanation examples for LM, MP and AVG given the recommended item *Catch me if you can* generated from the individual TF. Items are denoted in bold and aspects in italic.

anymore. Examples of user-targeted explanation generated using this approach can be found in Table 2.

3.3.4 Tag-cloud Explanations. Group recommendation explanations can also be visualized by mean of a Tag-cloud representation. Bilgic and Mooney [5] demonstrate how keyword-style explanations might improve recommendations’ perceived credibility and transparency. Fig. 2 shows an example of a tag-cloud representation in which we represent different aspects involved in the decision-making process of the predicted item. Therefore tags in our scheme are aspects, with the title of the wordcloud being the recommended item. We can encode visuals in terms of shape or colors based on how liked an aspect is by certain users of the group. In our case, colors are based on the predicted rating of the aspect (green being ratings close to 5, red close to 1). The height of aspects can be linked to its importance in the decision-making process, in our case they depend on the predicted rating. Our generated wordcloud are made for both the group as a whole as well as for the individual users. In the example in Fig. 2, only aspects with high predicted ratings are displayed.

4 CONCLUSIONS AND FUTURE WORKS

This paper presents a novel approach for providing explainable group recommendations, with a proof of concept of the use of an argumentative framework for generating group recommendations and explanations. The proposed approach is based on the concept of using an aspect-item framework as the foundation for



Figure 2: Wordcloud example for the movie *Catch me if you can*

an argumentation framework, which aims to help group members understand and accept the group recommendations. The underlying hybrid RS combines collaborative filtering and content-based recommendations, and is extended using some of the most used aggregation methods: Average satisfaction, Least Misery, and Most Pleasure. Furthermore, we explore several strategies to provide explanations, making use of the tripolar argumentation frameworks (TF) generated for the group and for each group member. More

specifically, we presented explanations for the whole group and explanations tailored for a specific group member, also considering the importance of trust, transparency, and privacy in explanations for GRS. Finally, we illustrated a very basic approach to provide graphical explanations in the form of tag-clouds.

Our approach presents several notable advantages. Its explanatory power may enhance users' understanding of the recommendations provided. Additionally, incorporating an argumentative framework enables the provision of both recommendations and explanations in a single approach. Furthermore, we propose a privacy-preserving version, ensuring the privacy of users of the group. Lastly, the inclusion of graphical explanations through tag clouds offers a different type of explanation. However, there are also some disadvantages to consider. The approach necessitates the creation of the A-I and TF components, which requires additional time and resources, together with domain-specific information necessary to connect aspects and items.

In future work, we plan to evaluate our approach, also to determine the "exact gain" of implementing our proposed system. First, we plan to perform an offline evaluation of the group recommendations by comparing them with baseline strategies. Such evaluation would be performed in different domains, and using both coupled and decoupled evaluation [1, 32]. Secondly, the explanations strategies needs to be evaluated. We plan to perform an online evaluation, through a set of user studies. We aim at evaluating the quality of explanations generated by the proposed approach for the entire group or a specific user in terms of transparency, effectiveness, trust, and privacy perception [26–29, 41], using social choice-based explanations as baseline [2, 43].

REFERENCES

- [1] Francesco Barile, Amra Delic, and Ladislav Peska. 2022. Tutorial on Offline Evaluation for Group Recommender Systems. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 702–705.
- [2] Francesco Barile, Shabnam Najafian, Tim Draws, Oana Inel, Alisa Rieger, Rishav Hada, and Nava Tintarev. 2021. Toward Benchmarking Group Explanations: Evaluating the Effect of Aggregation Strategies versus Explanation. In *Proceedings of the Perspectives on the Evaluation of Recommender Systems Workshop (PERSPECTIVES 2021)*.
- [3] Pietro Baroni, Giulia Comini, Antonio Rago, and Francesca Toni. 2017. Abstract games of argumentation strategy and game-theoretical argument strength. In *PRIMA 2017: Principles and Practice of Multi-Agent Systems: 20th International Conference, Nice, France, October 30–November 3, 2017, Proceedings* 20. Springer, 403–419.
- [4] Punam Bedi and Pooja Vashisth. 2015. Argumentation-enabled interest-based personalised recommender system. *Journal of Experimental & Theoretical Artificial Intelligence* 27, 2 (2015), 199–226.
- [5] Mustafa Bilgic and Raymond Mooney. 2005. Explaining Recommendations: Satisfaction vs. Promotion.
- [6] Cristian E Briguez, Maximiliano CD Budan, Cristhian AD Deagustini, Ana G Maguitman, Marcela Capobianco, and Guillermo R Simari. 2014. Argument-based mixed recommenders and their application to movie suggestion. *Expert Systems with Applications* 41, 14 (2014), 6467–6482.
- [7] Da Cao, Xiangnan He, Lianhai Miao, Guangyi Xiao, Hao Chen, and Jiao Xu. 2019. Social-enhanced attentive group recommendation. *IEEE Transactions on Knowledge and Data Engineering* 33, 3 (2019), 1195–1209.
- [8] Claudette Cayrol and Marie-Christine Lagasque-Schiex. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty: 8th European Conference, ECSQARU 2005, Barcelona, Spain, July 6–8, 2005. Proceedings* 8. Springer, 378–389.
- [9] CarlosIván Chesñevar, Ana Gabriela Maguitman, and María Paula González. 2009. Empowering recommendation technologies through argumentation. *Argumentation in artificial intelligence* (2009), 403–422.
- [10] Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence* 77, 2 (1995), 321–357.
- [11] Alexander Felfernig, Müslüm Atas, Thi Ngoc Trang Tran, Martin Stettinger, Seda Polat Erdeniz, and Gerhard Leitner. 2017. An Analysis of Group Recommendation Heuristics for High-and Low-Involvement Items. In *Advances in Artificial Intelligence: From Theory to Practice. IEA/AIE 2017 (Lecture Notes in Computer Science)*. Springer, 335–344. https://doi.org/10.1007/978-3-319-60042-0_39 International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2017 ; Conference date: 27-06-2017 Through 30-06-2017.
- [12] Alexander Felfernig, Ludovico Boratto, Martin Stettinger, and Marko Tkalcic. 2018. Explanations for Groups. In *Group Recommender Systems*. Springer, 105–126.
- [13] Alexander Felfernig, Nava Tintarev, Thi Ngoc Trang Tran, and Martin Stettinger. 2021. Designing Explanations for Group Recommender Systems.
- [14] Dov M Gabbay. 2016. Logical foundations for bipolar and tripolar argumentation networks: preliminary results. *Journal of Logic and Computation* 26, 1 (2016), 247–292.
- [15] Lewis R Goldberg. 1993. The structure of phenotypic personality traits. *American psychologist* 48, 1 (1993), 26.
- [16] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. 2010. *Recommender systems: an introduction*. Cambridge University Press.
- [17] Yogesh Jhamb and Yi Fang. 2017. A dual-perspective latent factor model for group-aware social event recommendation. *Information Processing & Management* 53, 3 (2017), 559–576.
- [18] Yucheng Jin, Karsten Seipp, Erik Duval, and Katrien Verbert. 2016. Go With the Flow: Effects of Transparency and User Control on Targeted Advertising Using Flow Charts. 68–75. <https://doi.org/10.1145/2909132.2909269>
- [19] Öykü Kapcak, Simone Spagnoli, Vincent Robbmond, Soumitri Vadali, Shabnam Najafian, and Nava Tintarev. 2018. TourExplain: A Crowdsourcing Pipeline for Generating Explanations for Groups of Tourists. In *RecTour@RecSys*.
- [20] Jerry S Kelly. 2013. *Social choice theory: An introduction*. Springer Science & Business Media.
- [21] Qianqiao Liang, Xiaolin Zheng, Yan Wang, and Mengying Zhu. 2021. O3ERS: An explainable recommendation system with online learning, online recommendation, and online explanation. *Information Sciences* 562 (2021), 94–115. <https://doi.org/10.1016/j.ins.2020.12.070>
- [22] Xin Luo, Mengchu Zhou, Yundi Xia, and Qingsheng Zhu. 2014. An efficient non-negative matrix-factorization-based approach to collaborative filtering for recommender systems. *IEEE Transactions on Industrial Informatics* 10, 2 (2014), 1273–1284.
- [23] Judith Masthoff. 2004. Group modeling: Selecting a sequence of television items to suit a group of viewers. In *Personalized digital television*. Springer, 93–141.
- [24] Judith Masthoff. 2015. Group recommender systems: aggregation, satisfaction and group attributes. In *recommender systems handbook*. Springer, 743–776.
- [25] Judith Masthoff and Amra Delic. 2022. Group Recommender Systems: Beyond Preference Aggregation. In *Recommender Systems Handbook*. Springer, 381–420.
- [26] Shabnam Najafian, Amra Delic, Marko Tkalcic, and Nava Tintarev. 2021. Factors Influencing Privacy Concern for Explanations of Group Recommendation. 14–23. <https://doi.org/10.1145/3450613.3456845>
- [27] Shabnam Najafian, Tim Draws, Francesco Barile, Marko Tkalcic, Jie Yang, and Nava Tintarev. 2021. Exploring User Concerns about Disclosing Location and Emotion Information in Group Recommendations. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*. 155–164.
- [28] Shabnam Najafian, Oana Inel, and Nava Tintarev. 2020. Someone really wanted that song but it was not me! Evaluating Which Information to Disclose in Explanations for Group Recommendations. In *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion*. 85–86.
- [29] Shabnam Najafian and Nava Tintarev. 2018. Generating Consensus Explanations for Group Recommendations: an exploratory study. 245–250. <https://doi.org/10.1145/3213586.3225231>
- [30] Sidra Naveed, Tim Donkers, and Jürgen Ziegler. 2018. Argumentation-based explanations in recommender systems: Conceptual framework and empirical results. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*. 293–298.
- [31] Thuy-Ngoc Nguyen and Francesco Ricci. 2018. A chat-based group recommender system for tourism. *Information Technology & Tourism* 18 (04 2018). <https://doi.org/10.1007/s40558-017-0099-y>
- [32] Ladislav Peska and Ladislav Malecek. 2021. Coupled or Decoupled Evaluation for Group Recommendation Methods?. In *Perspectives@ RecSys*.
- [33] Pearl Pu, Li Chen, and Rong Hu. 2012. Evaluating recommender systems from the user's perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction* 22 (2012), 317 – 355.
- [34] Lara Quijano-Sanchez, Christian Sauer, Juan Recio-García, and Belén Díaz-Agudo. 2017. Make it Personal: A Social Explanation System Applied to Group Recommendations. *Expert Systems with Applications* 76 (01 2017). <https://doi.org/10.1016/j.eswa.2017.01.045>
- [35] Antonio Rago, Oana Cocarascu, Christos Bechlivanidis, David Lagnado, and Francesca Toni. 2021. Argumentative explanations for interactive recommendations. *Artificial Intelligence* 296 (2021), 103506. <https://doi.org/10.1016/j.artint.2021.103506>

- 2021.103506
- [36] Christophe Senot, Dimitre Kostadinov, Makram Bouzid, Jérôme Picault, Armen Aghasaryan, and Cédric Bernier. 2010. Analysis of strategies for building group profiles. In *International Conference on User Modeling, Adaptation, and Personalization*. Springer, 40–51.
 - [37] Rashmi Sinha and Kirsten Medhurst. 2002. The role of transparency in recommender systems. 830. <https://doi.org/10.1145/506558.506619>
 - [38] Panagiotis Symeonidis, Alexandros Nanopoulos, and Yannis Manolopoulos. 2008. Providing Justifications in Recommender Systems. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 38, 6 (2008), 1262–1272. <https://doi.org/10.1109/TSMCA.2008.2003969>
 - [39] Juan CL Teze, Sebastian Gottifredi, Alejandro J García, and Guillermo R Simari. 2015. Improving argumentation-based recommender systems through context-adaptable selection criteria. *Expert Systems with Applications* 42, 21 (2015), 8243–8258.
 - [40] Nava Tintarev and Judith Masthoff. 2015. *Explaining Recommendations: Design and Evaluation*. Springer Science+Business Media, 353–382. https://doi.org/10.1007/978-1-4899-7637-6_10
 - [41] Nava Tintarev and Judith Masthoff. 2022. Beyond explaining single item recommendations. In *Recommender Systems Handbook*. Springer, 711–756.
 - [42] St Toulmin. 1958. *The Uses of Argument* Cambridge University Press. Cambridge, UK (1958).
 - [43] Thi Ngoc Trang Tran, Müslüm Atas, Alexander Felfernig, Viet Man Le, Ralph Samer, and Martin Stettinger. 2019. Towards social choice-based explanations in group recommender systems. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*. 13–21.
 - [44] Katrien Verbert, Denis Parra, Peter Brusilovsky, and erik duval. 2013. Visualizing recommendations to support exploration, transparency and controllability. *International Conference on Intelligent User Interfaces, Proceedings IUI*, 351–362. <https://doi.org/10.1145/2449396.2449442>
 - [45] Jesse Vig, Shilad Sen, and John Riedl. 2009. Tagsplanations: explaining recommendations using tags. 47–56. <https://doi.org/10.1145/1502650.1502661>
 - [46] Wei Wang, Guangquan Zhang, and Jie Lu. 2017. Hierarchy visualization for group recommender systems. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 49, 6 (2017), 1152–1163.
 - [47] Yongfeng Zhang and Xu Chen. 2020. Explainable Recommendation: A Survey and New Perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101. <https://doi.org/10.1561/15000000066>