Towards LLM-Enhanced Group Recommender Systems

Sebastian **Lubos**[1,*], Thi Ngoc Trang **Tran**[1], Viet-Man **Le**[1], Damian **Garber**[1], Manuel **Henrich**[2], Reinhard **Willfort**[3] and Jeremias **Fuchs**[4]

[1]*Graz University of Technology, Inffeldgasse 16b, Graz, 8010, Austria*

[2]*Uniquare Software Development, Lannerweg 9, Krumpendorf, 9201, Austria*

[3]*Innovation Service Network, Grabenstraße 231, Graz, 8045, Austria*

[4]*Morgendigital, Leopoldstraße 20, Innsbruck, 6020, Austria*

Abstract

In contrast to single-user recommender systems, group recommender systems are designed to generate and explain recommendations for groups. This group-oriented setting introduces additional complexities, as several factors—absent in individual contexts—must be addressed. These include understanding group dynamics (e.g., social dependencies within the group), defining effective decision-making processes, ensuring that recommendations are suitable for all group members, and providing group-level explanations as well as explanations for individual users. In this paper, we analyze in which way large language models (LLMs) can support these aspects and help to increase the overall decision support quality and applicability of group recommender systems.

Keywords

Group Recommender Systems, Large Language Models, Decision Making in Groups

1. Introduction

Nowadays, recommender systems are a foundational component of many digital platforms, which help to personalize item suggestions [1]. Traditionally, recommender systems analyze the preferences and behaviors of *single users* to optimize item rankings. In contrast, *group recommender systems* [2, 3, 4] support the decision-making of multiple users by aggregating their individual preferences. This approach triggers various challenges such as the resolution of conflicting preferences, the assurance of fairness with regard to the determined recommendations, and the identification of appropriate decision (aggregation) strategies. While recommender systems for single users focus on personalization, group recommender systems must also take into account the aspects of negotiation (among users) and compromise with regard to the proposed items.

The application of group recommender systems is relevant in the context of decision tasks with a focus on *collaborative decision-making* [5]. For example, families or groups of friends commonly make joint decisions about holiday destinations, accommodations, and itineraries. In a similar fashion, restaurant selection is often performed by groups (e.g., restaurant selection for a Christmas party). The group decision support tool DOODLE[1] can be regarded as a very basic form of group recommender system where preferred time slots can be interpreted as implicit recommendation (based on majority voting) [6]. Although entertainment environments such as Netflix[2] have already performed initial studies regarding the applicability of group recommender systems [7], their productive platforms are still focused on

[1]https://doodle.com/

[2]https://netflix.com

individual users, i.e., not groups. Music recommendation provides another relevant application domain for group recommender systems, for example, in terms of shared listening environments for parties or gyms. However, popular platforms like Spotify[3] have not integrated group recommendation with the exception of basic collaborative playlists [8]. In software engineering, group recommender systems are applied specifically in the context of requirements engineering scenarios where the preferences and background knowledge of stakeholders are an input for recommender systems that support the prioritization of software requirements [9]. For an overview of applications of group recommender systems, we refer to Felfernig et al. [4].

Despite the existence of different promising application areas for group recommender systems, these systems are not widely applied and accepted. On the one hand, a reason behind is a *lack of flexibility* in the support of group decision processes, for example, in taking over a more supportive role in contrast to the proposal of concrete recommendations [10]. On the other hand, group decisions are often based on *sensitive or private information* that people do not want to share with others in an explicit fashion. Examples thereof are physical conditions in holiday round-trip planning, interpersonal likings in group formation or personal selection, or personal preferences regarding movies [11]. Furthermore, some group decision scenarios suffer from *hidden agendas*, where individuals strategically withhold or distort their preferences to manipulate the group outcome [12]. Examples thereof can be found in professional environments, for example, strategy meetings, recruitment sessions, or funding decisions. In contrast, domains involving leisure and less-sensitive content, such as restaurants and meeting dates, show a higher acceptance of group decision support.

In addition to technical, privacy, and manipulation issues, *cultural and societal values* could also play a role in the adoption of group recommender systems [13, 14]. In Western societies, we can observe a tendency towards individualism, which can translate to a strong focus on individual preferences. This can also counteract the acceptance of group recommender systems, which require, to some extent, a preparedness to compromise and openness. In contrast, in more collectivist cultures, group harmony is prioritized higher, which could increase an interest in group recommender systems. As a consequence, sociocultural aspects have to be regarded as acceptance factors besides technical performance.

Finally, the internal structure of the group (i.e., homogeneous vs. heterogeneous) has an impact on group recommender acceptance. In homogeneous groups such as close friends with shared tastes, basic preference aggregation could be an accepted recommendation mechanism. Heterogeneous groups such as large (often cross-cultural) software development teams might be confronted with diverse and often conflicting preferences. In such scenarios, simple aggregation functions reach their limits, and more sophisticated decision support is needed that considers aspects such as fairness, the need for negotiating decisions, and resolving conflicts. In this paper, we focus on the aspect of making the decision support provided by group recommender systems more flexible. In the line of Lin et al. [15] and Zhang et al. [16], we analyze *different ways of how large language models (LLMs) can support group decision scenarios*.

The remainder of this paper is organized as follows. Section 2 provides an overview of basic algorithms for group recommender systems and highlights how LLMs can enhance these algorithms. In Section 3, we discuss existing preference elicitation methods for group recommenders and how these methods can be improved on the basis of LLMs. Section 4 outlines how LLMs can support explanations in the context of group recommendation scenarios. Section 5 shows how synergy effects can be achieved by combining LLMs with insights from psychological theories of decision making. Section 6 summarizes open research challenges related to the integration of LLMs into group recommender systems. The paper is concluded in Section 7.

2. Algorithms for Group Recommender Systems

Group-based collaborative filtering (GCF) extends basic collaborative filtering for single users by aggregating the preferences of single users into a group model.[4] Examples of such aggregation strategies are

---

[3]https://www.spotify.com

[4]For a detailed discussion of different related group recommendation approaches, we refer to [4, 17, 18].

the *average* strategy (the mean of all user preferences is used as an estimate for the degree of group acceptance), *least misery* (the satisfaction level / item acceptance of the user with the lowest rating is used as an estimate for the group acceptance level), and most pleasure (the satisfaction level / item acceptance of the user with the highest rating is used as an estimate for the group acceptance level). Consider a group of three users rating four movies as shown in Table 1. Using *least misery*, the group score for Movie A would be $\min(4, 5, 2) = 2$, while the average strategy would result in $\frac{4+5+2}{3} = 3.67$. Such aggregation functions often oversimplify group dynamics and the history of the decision process.

Table 1
Example of group member (user) ratings for four movies.

User	Movie A	Movie B	Movie C	Movie D
User 1	4	3	5	2
User 2	5	4	4	1
User 3	2	4	5	3

Content-based filtering for group recommendation (CBFG) is based on the idea of constructing and aggregating user profiles [19], for example, as vectors of preferred item features. If User X likes movies with the features "romantic" and "comedy" and User Y prefers "action" and "thriller", a combined profile must balance these dimensions. Let the profiles of two users be $\vec{u}_X = [1, 0.8, 0, 0]$ and $\vec{u}_Y = [0, 0, 1, 0.9]$ for the genres [romance, comedy, action, thriller]. A simple average profile is $[0.5, 0.4, 0.5, 0.45]$.

Critiquing-based group recommendation (CRITG) – in contrast to GCF and CBFG – supports user feedback cycles in terms of critiques [20]. Examples of critiques in restaurant recommendation are "less expensive", "nearer", or "higher food quality". Such critiques are aggregated where conflict resolution can become an issue due to conflicting critiques provided by different users (e.g., "high quality food" vs. "less expensive" could lead to a situation where no recommendation can be identified) [4].

LLM-related Potentials. LLMs have the potential to significantly enhance algorithmic approaches for group recommender systems by interpreting user input in a more human-centered way. In the context of GCF, LLMs can move beyond numerical ratings by *interpreting free-form feedback* such as "I loved the plot but hated the violence" and *transforming it into preference vectors*. In CBFG, LLMs can *extract item attributes* from unstructured sources with the goal of enriching user and item profiles. In the context of CRITG, LLMs can *mediate conflicting critiques* using generative dialog strategies that *suggest reformulations*. Most importantly, LLMs can be exploited to dynamically *propose appropriate preference aggregation and decision strategies*. For example, on the basis of textual user preference definitions such as "User 2: I really would like to see Movie C but I am definitely not interested in watching Movie D", "User 1: I am completely flexible regarding the movie we intend to watch", and "User 3: I am not in the mood for horror and action movies, the only options for today would be Movie B or Movie C", an LLM-based group recommender could flexibly infer that *Movie C* is an acceptable option for all users. However, LLM-enhanced group recommendation could also support open recommendation processes, for example, by flexibly extending the choice set with a new option "Movie E" which might be even better compared to "Movie C". Having available the dialog history of the users (could be regarded as a kind of retrieval-augmented generation), an LLM-based group recommender could also adapt the used preference aggregation strategy. For example, if a dialog analysis comes to the conclusion that over a longer time period the preferences of User 3 have not been taken into account, the LLM could choose a corresponding fairness-aware recommendation strategy (corresponding instructions should be included in the LLM prompting).

3. Preference Elicitation in Group Recommender Systems

In the context of group recommender systems, the complexity of preference elicitation increases compared to single-user recommender systems due to the presence of different and possibly conflicting preferences. In group decision contexts, preferences can be elicited by collecting explicit ratings

from users (i.e., group members), on the basis of user-individual item rankings, on the basis of item comparisons, but also in an implicit fashion, for example, by observing user interaction patterns when interacting with the recommender system [21]. In the following, the elicited user preferences are ranked, for example, with aggregation functions (e.g., least misery or average). Such a way of preference elicitation can work well in standard settings but can be suboptimal in informal group decision settings where users prefer to articulate their preferences and feedback in an informal fashion, i.e., not in terms of ratings. More advanced elicitation approaches leverage conversational interfaces to capture individual inputs with higher engagement, for example, user preferences can be collected in terms of a *pro/con analysis* which is then used as an input for the determination of a group recommendation [22]. However, challenges still remain in mediating conflicts and interpreting potentially ambiguous inputs.

LLM-related Potentials. LLMs offer the potential of providing a more natural and context-rich way of *interactive preference elicitation.* On the basis of LLMs, basic item ratings can be replaced by *free-form conversation* with the possibility of extracting *more nuanced user preferences* [23]. For example, for informally articulated user preferences regarding a movie to watch, an LLM prompt "What kind of movie would everyone enjoy tonight?" could provide a sentiment estimate, detect conflicts, and propose *explanations in real-time.* Beyond existing group recommendation approaches, LLMs can help to detect disengaged users and try to involve them in the group decision process. Furthermore, LLMs can support *conflict resolution*, for example, by engaging users in further discussions with the goal of finding a compromise. Such preference elicitation interfaces can be *multi-modal*, i.e., combining text-based preferences with information from video streams and audio tracks from the current session [24]. For example, in the context of group-based requirements prioritization [9], two stakeholders could disagree regarding the technical risk of a specific requirement A. By analyzing related dialogs, an LLM could detect such contradictions and propose countermeasures ranging from contacting the affected stakeholders and stimulating discussions but also to provide more in-depth information which could even contribute to immediately resolve the contradiction. For example, an LLM could provide the following (local) information to the stakeholders: "a similar requirement has been implemented for customer Y – related technical risks can be neglected".

4. Explanations in Group Recommender Systems

Major objectives of providing explanations in recommender systems are to improve *transparency* (i.e., to help users understand the system output in terms of *why?* explanations) and create *trust* [25]. In group recommender systems, explanations play an even more important role, since explanations need to help mediate between diverging user preferences and justify tradeoffs made during the phase of preference aggregation. Explanations in the group recommendation context can be based on *scores* (e.g., this restaurant is recommended since it has the highest average rating). Following this approach, explanations are easy to understand. However, major properties such as conflicts, group dynamics, and compromises are not taken into account. Following a *feature-based* explanation approach, features are used to explain a recommendation. A related example is: "*This movie is recommended since it satisfies at least two members' preferences for comedy and one member's preference for history.*"

An important aspect in explanation generation is to present explanations *without revealing sensitive user preferences.* Obfuscating preferences can lead to lower clarity of recommendations, but support privacy. The contents of Table 2 can be used for explanation purposes by offering an aggregated view on the preferences of the group and thus respecting user boundaries by abstracting preference information to avoid personal attribution. Such explanations are also denoted as *group-level explanations*, which reflect the preference landscape of the whole group. In contrast, *personalized explanations* refer to the preferences of a single group member and are often not shown to the other group members. A challenge in explanation generation is to find a balance between group-level and personalized explanations.

LLM-related Potentials. LLMs can generate context-sensitive and *more naturalistic explanations* [26]. In group decision contexts, specific aspects of the decision process have to be taken into account, for example, the sentiments of different group members, potential compromises in conflicting situations, and

Table 2

Aggregated feature weights as a basis for group-level explanations.

Feature	Average Group Weight	Impact on Recommendation
Action Genre	0.75	High
Romance Genre	0.40	Medium
Comedy Genre	0.60	High
Historical Setting	0.30	Low

social relationships between group members. Due to their generative capabilities, LLMs can *adapt the granularity* of an explanation based on different user roles in a decision process. For example, a detailed explanation is provided to engaged users with a corresponding topic-wise background knowledge, whereas only a summary is provided for other users with low related expertise. Explanations can also be provided for resolved conflicts and corresponding compromises, for example, "*since two members prefer hiking and one prefers city tours, we recommend a trip that includes both activities.*"

5. Psychological Decision Models in Group Recommender Systems

Psychological models of human decision-making provide insights into the patterns groups typically follow. Taking into account such models can help to significantly enhance the effectiveness and quality of decision support [27, 28, 29, 30, 31, 32]. A related example concept is denoted as *emotional contagion* [33], which represents the idea that the mood of one group member can influence the mood of other group members. Consequently, if a specific group member is very enthusiastic about an alternative (e.g., a tourist destination to visit), this could have an impact on other group members in such a way that they increase their evaluation of the mentioned alternative. If group recommender systems have knowledge about the social relationships between different users, this information can be exploited by automatically updating the corresponding preferences. Such social relationships can be represented in terms of *emotional alignment scores* [34] as shown in Table 3. Application domains where such an emotional resonance plays an important role are entertainment (e.g., which movie to watch?) and restaurant decisions – in such contexts, emotional resonance can affect group satisfaction [35].

Table 3

Emotional alignment scores among group members with a strong alignment between *User 1* and *User 2*.

User	User 1	User 2	User 3
User 1	1.00	0.85	0.60
User 2	0.85	1.00	0.50
User 3	0.60	0.50	1.00

Table 4

Preference indicators from multimodal video data.

Modality	Feature	Preference Indicator
Facial Expression	Smiling duration	Positive sentiment
Eye Gaze	Focus on screen/object	Attention/interest
Voice Tone	Pitch variation	Excitement/frustration
Body Posture	Leaning forward	Engagement

Another related concept denoted as *groupthink* [36] represents situations where the desire for harmony or conformity has a negative impact on the quality of a decision outcome. In the context of group recommender systems, such situations can occur if dominant group members/users (e.g., leaders on the basis of their role in the group) implicitly override the real preferences of other users. In

online recommendation, groupthink might not be detected if only basic preference data is available – a recommender system might infer group consensus. To address such issues, a recommender needs to be aware of user interaction patterns. For example, users always follow the proposed rating of a specific other person, or users rarely give feedback (or only positive feedback) on the evaluations of other users.

Group polarization is the effect that groups as a whole take more extreme decisions compared to their individual preferences [27]. For example, if individual group members have an interest in thriller movies, a group as a whole might choose a horror movie that is in some sense more extreme than the individual group members' preferences. A similar situation might occur in the context of credit decisions where the preparedness to take risks of a whole group exceeds the preparedness to take risks of individual group members [27]. Consequently, it is important to be able to identify such decision patterns, which can be done, for example, by analyzing preference changes over time by individual group members and the divergence between the final decision and the initially articulated preferences. With such an analysis, an offering of more moderate compromise options can help to counteract polarization.

LLM-related Potentials. LLMs can help to integrate psychological decision models more deeply into group recommender systems by *interpreting the subtle signals embedded in natural language and dialog structure*. In this context, emotional contagion can be detected, for example, by analyzing the *sentiments in group chats* of individual group members throughout the decision process. In a similar fashion, LLMs can help to infer groupthink by *analyzing diversity or dissent in the opinions of group members*. To counteract groupthink, an LLM can generate alternative options (not considered up to now) and also evaluate existing options differently. Given a chat history or other forms of interaction patterns (e.g., a real-time video sequence of a group decision session [24]), group members can be analyzed with regard to their role in the decision process, i.e., who is the moderator, who provides ideas, who initiates detailed discussions, and who remains silent in most of the cases. Such information can be exploited for behavioral modeling and determining countermeasures that help to avoid decision biases [24]. This can be regarded as a kind of failure mode in group decision scenarios which activates relevant countermeasures.

6. Research Issues in LLM Integration

The integration of LLMs into group recommender systems triggers a couple of research issues.

Recommendation of Decision Strategies. Existing preference aggregation strategies are not appropriate in every case and there does not exist a general set of rules that clearly specify when to apply which strategy [5]. LLMs can help out by analyzing the decision context and the preferences/roles of the individual group members. With this information, LLMs can propose corresponding decision strategies. A related research challenge is the automated design, validation, and explanation of proposed decision strategies (also for increasing trust in strategy recommendations). Note that decision strategies could be complex, for example, due to the unclear employment status of a family member, an apartment purchasing decision needs to be deferred. Furthermore, investment decisions need to be optimized, which may require integrating LLM-based decision support with corresponding financial calculators to ensure the correctness of calculations

Assuring Fairness in LLM-based Recommendations. On the basis of diverse and potentially contradicting preferences of individual group members, LLMs might have a tendency to over-represent the more dominant opinions in a group decision setting. To counteract such situations (e.g., in the context of decision scenarios with culturally diverse groups), corresponding guardrails need to be established that ensure that specific compliance policies are taken into account. For example, to assure fairness in group decisions, bias mitigation mechanisms have to be integrated into LLM architectures and/or corresponding pre- and post-processing components [37].

Supporting Realtime Group Modeling. LLMs offer the capability of so-called adaptive group modeling, where new statements of individual group members can lead to a new interpretation of the overall status of a group decision process. For example, real-time video-based group decision making [24] can profit from LLM capabilities to analyze the current status of the decision process by immediately providing updates on the basis of the interactions of group members. If one group member proposes a new decision alternative and provides an argument for or against a specific alternative, an LLM-based group recommender could immediately react in terms of indicating similar alternatives proposed in the past with arguments that led to rejection.

Multi-modal Preference Elicitation. Many group recommender systems are based on the idea of collecting preferences on the basis of simple text-based (e.g., critiquing) or numeric user feedback [4]. However, in group decision scenarios, there often exist different sources for the identification of user preferences [38]. Preferences can be inferred from direct user feedback, as well as from visual, audio, and behavioral data (e.g., user navigation behavior). For example, take a look at the example depicted in Table 4 that focuses on preference elicitation from video streams. These basic indicators of user behavior can be interpreted by an LLM for the purpose of identifying the next relevant tasks to be performed in the group decision process.

Ethical and Privacy Issues. Preference extraction from multiple data sources, specifically from video feeds and personal chats, requires the inclusion of privacy-preserving mechanisms that proactively help to avoid an unintended "transfer" of sensitive data [39]. Furthermore, fine-grained user consent models and intuitive preference management interfaces need to be provided for group decision scenarios.

7. Conclusions

Group recommender systems provide decision support features in domains such as entertainment and software engineering. However, the application of group recommenders is still underrepresented in commercial contexts for various reasons, ranging from limited technological support, the need for sharing potentially sensitive data, and the danger of manipulated decision outcomes triggered by hidden agendas. This paper includes a short overview of the state of the art in group recommendation with a specific focus on the potential of integrating LLMs into group recommender system-related algorithms, preference elicitation mechanisms, explanations, and psychological models.

8. Acknowledgements

Declaration on Generative AI

During the preparation of this work, the authors used *ChatGPT* and *Grammarly* in order to: *Grammar and spelling check, Paraphrase and reword.* After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] D. Jannach, M. Zanker, A. Felfernig, G. Friedrich, Recommender Systems: An Introduction, 1st ed., Cambridge University Press, USA, 2010.

[2] A. Jameson, B. Smyth, Recommendation to Groups, Springer, Berlin, Heidelberg, 2007, pp. 596–627. doi:10.1007/978-3-540-72079-9_20.

[3] J. Masthoff, Group recommender systems: Combining individual models, in: F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), Recommender Systems Handbook, Springer US, Boston, MA, 2011, pp. 677–702. doi:10.1007/978-0-387-85820-3_21.

[4] A. Felfernig, L. Boratto, M. Stettinger, M. Tkalčič, Group Recommender Systems: An Introduction, Springer, 2024.

[5] A. Jameson, M. Willemsen, A. Felfernig, Individual and group decision making and recommender systems, in: F. Ricci, L. Rokach, B. Shapira (Eds.), Recommender Systems Handbook, Springer, New York, NY, 2022, pp. 789–832. doi:10.1007/978-1-0716-2197-4_21.

[6] K. Reinecke, M. K. Nguyen, A. Bernstein, M. Näf, K. Z. Gajos, Doodle around the world: online scheduling behavior reflects cultural differences in time perception and group decision-making, in: Conference on Computer Supported Cooperative Work, CSCW '13, ACM, New York, NY, USA, 2013, pp. 45–54. doi:10.1145/2441776.2441784.

[7] S. Berry, S. Fazzio, Y. Zhou, B. Scott, L. Francisco-Revilla, Netflix recommendations for groups, in: 73rd ASIS&T Annual Meeting on Navigating Streams in an Information Ecosystem - Volume 47, ASIS&T '10, American Society for Information Science, USA, 2010.

[8] D. Kwak, S. Park, I. Cha, H. Kim, Y.-K. Lim, Investigating the potential of group recommendation systems as a medium of social interactions: A case of spotify blend experiences between two users, in: CHI '24, ACM, New York, NY, USA, 2024. doi:10.1145/3613904.3642544.

[9] G. Ninaus, A. Felfernig, M. Stettinger, S. Reiterer, G. Leitner, L. Weninger, W. Schanil, IntelliReq: Intelligent Techniques for Software Requirements Engineering, ECAI'14, 2014, pp. 1161–1166.

[10] D. Jannach, A. Delić, F. Ricci, M. Zanker, Rethinking Group Recommender Systems in the Era of Generative AI: From One-Shot Recommendations to Agentic Group Decision Support, 2025. URL: https://arxiv.org/abs/2507.00535. arXiv:2507.00535.

[11] Z. Luo, Z. Chen, A privacy preserving group recommender based on cooperative perturbation, in: International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery, 2014, pp. 106–111. doi:10.1109/CyberC.2014.26.

[12] T. N. T. Tran, A. Felfernig, V. M. Le, M. Atas, M. Stettinger, R. Samer, User interfaces for counter-acting decision manipulation in group recommender systems, in: Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization, UMAP'19 Adjunct, ACM, New York, NY, USA, 2019, pp. 93–98. doi:10.1145/3314183.3324977.

[13] G. Hofstede, Culture's consequences: Comparing values, behaviors, institutions and organizations across nations, Sage Publications, 2001.

[14] D. Zhang, P. Lowry, L. Zhou, X. Fu, The impact of individualism-collectivism, social presence, and group diversity on group decision making under majority influence, J. Manage. Inf. Syst. 23 (2007) 53–80. doi:10.2753/MIS0742-1222230404.

[15] J. Lin, X. Dai, Y. Xi, W. Liu, B. Chen, H. Zhang, Y. Liu, C. Wu, X. Li, C. Zhu, H. Guo, Y. Yu, R. Tang, W. Zhang, How can recommender systems benefit from large language models: A survey, ACM Trans. Inf. Syst. 43 (2025). doi:10.1145/3678004.

[16] J. Zhang, K. Bao, Y. Zhang, W. Wang, F. Feng, X. He, Large Language Models for Recommendation: Progresses and Future Directions, in: Companion Proceedings of the ACM Web Conference 2024, WWW '24, ACM, New York, NY, USA, 2024, pp. 1268–1271. doi:10.1145/3589335.3641247.

[17] L. Boratto, A. Felfernig, Group recommendations, in: Collaborative Recommendations, World Scientific, 2019, pp. 203–232. doi:10.1142/9789813275355_0006.

[18] J. Masthoff, A. Delić, Group recommender systems: Beyond preference aggregation, in: F. Ricci, L. Rokach, B. Shapira (Eds.), Recommender Systems Handbook, Springer US, New York, NY, 2022, pp. 381–420. doi:10.1007/978-1-0716-2197-4_10.

[19] Y. Pérez-Almaguer, R. Yera, A. A. Alzahrani, L. Martínez, Content-based group recommender systems: A general taxonomy and further improvements, Expert Systems with Applications 184 (2021) 115444. doi:https://doi.org/10.1016/j.eswa.2021.115444.

[20] K. McCarthy, M. Salamó, L. Coyle, L. McGinty, B. Smyth, P. Nixon, Group recommender systems: a critiquing based approach, in: 11th International Conference on Intelligent User Interfaces, IUI '06, ACM, New York, NY, USA, 2006, pp. 267–269. doi:10.1145/1111449.1111506.

[21] I. Garcia, S. Pajares, L. Sebastia, E. Onaindia, Preference elicitation techniques for group recommender systems, Information Sciences 189 (2012) 155–175. doi:`10.1016/j.ins.2011.11.037`.

[22] A. Felfernig, T. Gruber, G. Brandner, P. Blazek, M. Stettinger, Customizing Events with EventHelpr, in: 8th Intl. Conf. on Mass Customization and Personalization, MCP-CE 2018, 2018, pp. 88–91.

[23] S. Feng, Z. Lang, J. He, H. Zhang, W. Chen, J. Cao, A group recommendation method based on automatically integrating members' preferences via taking advantages of llm, Information Sciences 709 (2025) 122067. doi:`https://doi.org/10.1016/j.ins.2025.122067`.

[24] S. Lubos, A. Felfernig, D. Garber, V.-M. Le, M. Henrich, R. Willfort, J. Fuchs, Towards Group Decision Support with LLM-based Meeting Analysis, in: 33rd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '25, ACM, 2025, pp. 331–335. doi:`10.1145/3708319.3733646`.

[25] N. Tintarev, J. Masthoff, Effective explanations of recommendations: user-centered design, in: ACM Conference on Recommender Systems, RecSys '07, ACM, New York, NY, USA, 2007, pp. 153–156. doi:`10.1145/1297231.1297259`.

[26] S. Lubos, T. N. T. Tran, A. Felfernig, S. Polat Erdeniz, V.-M. Le, LLM-generated Explanations for Recommender Systems, in: 32nd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '24, ACM, New York, NY, USA, 2024, pp. 276–285. doi:`10.1145/3631700.3665185`.

[27] M. Atas, S. Reiterer, A. Felfernig, T. N. T. Tran, M. Stettinger, Polarization effects in group decisions, in: Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, UMAP '18, ACM, New York, NY, USA, 2018, pp. 305–310. doi:`10.1145/3213586.3225242`.

[28] M. Atas, R. Samer, A. Felfernig, T. N. T. Tran, S. P. Erdeniz, M. Stettinger, Socially-aware diagnosis for constraint-based recommendation, in: 27th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '19, ACM, New York, NY, USA, 2019, pp. 121–129. doi:`10.1145/3320435.3320436`.

[29] P. Bota, J. Brito, A. Fred, P. Cesar, H. Silva, A real-world dataset of group emotion experiences based on physiological data, Sci Data 11 (2024). doi:`10.1145/1297231.1297259`.

[30] L. Chen, M. de Gemmis, A. Felfernig, P. Lops, F. Ricci, G. Semeraro, Human decision making and recommender systems, ACM Trans. Int. Intell. S. 3 (2013). doi:`10.1145/2533670.2533675`.

[31] D. Contreras, M. Salamó, L. Boratto, Integrating Collaboration and Leadership in Conversational Group Recommender Systems, ACM Trans. Inf. Syst. 39 (2021). doi:`10.1145/3462759`.

[32] M. Stettinger, A. Felfernig, G. Leitner, S. Reiterer, M. Jeran, Counteracting serial position effects in the choicla group decision support environment, in: 20th International Conference on Intelligent User Interfaces, IUI '15, ACM, New York, NY, USA, 2015, pp. 148–157. doi:`10.1145/2678025.2701391`.

[33] J. Masthoff, A. Gatt, In pursuit of satisfaction and the prevention of embarrassment: affective state in group recommender systems 16 (2006) 281–319. doi:`10.1007/s11257-006-9008-3`.

[34] X. Zhang, R. Xie, Y. Lyu, X. Xin, P. Ren, M. Liang, B. Zhang, Z. Kang, M. de Rijke, Z. Ren, Towards empathetic conversational recommender systems, RecSys '24, ACM, New York, NY, USA, 2024, pp. 84–93. doi:`10.1145/3640457.3688133`.

[35] T. N. T. Tran, A. Felfernig, N. Tintarev, Humanized recommender systems: State-of-the-art and research issues, ACM Trans. Int. Intell. S. 11 (2021). doi:`10.1145/3446906`.

[36] J. K. Esser, Alive and Well after 25 Years: A Review of Groupthink Research, Organizational Behavior and Human Decision Processes 73 (1998) 116–141. doi:`10.1006/obhd.1998.2758`.

[37] A. Tommasel, Fairness Matters: A look at LLM-generated group recommendations, in: 18th ACM Conference on Recommender Systems, RecSys '24, ACM, New York, NY, USA, 2024, pp. 993–998. doi:`10.1145/3640457.3688182`.

[38] O. Alvarado, N. N. Htun, Y. Jin, K. Verbert, A Systematic Review of Interaction Design Strategies for Group Recommendation Systems, ACM Hum.-Comput. Interact. 6 (2022). doi:`10.1145/3555161`.

[39] Y. Ge, S. Liu, Z. Fu, J. Tan, Z. Li, S. Xu, Y. Li, Y. Xian, Y. Zhang, A Survey on Trustworthy Recommender Systems, ACM Trans. Recomm. Syst. 3 (2024). doi:`10.1145/3652891`.