

Concentrating on the Impact: Consequence-based Explanations in Recommender Systems

Sebastian Lubos¹, Thi Ngoc Trang Tran¹, Seda Polat Erdeniz¹, Merfat El Mansi¹, Alexander Felfernig¹, Manfred Wundara² and Gerhard Leitner³

¹*Institute of Software Technology, Graz University of Technology, Inffeldgasse 16b, 8010 Graz, Austria*

²*Magistrat Villach, Rathausplatz 1, 9500 Villach, Austria*

³*Interactive Systems Group, University of Klagenfurt, Universitätsstrasse 65-67, 9020 Klagenfurt am Wörthersee, Austria*

Abstract

Recommender systems assist users in decision-making, where the presentation of recommended items and their explanations are critical factors for enhancing the overall user experience. Although various methods for generating explanations have been proposed, there is still room for improvement, particularly for users who lack expertise in a specific item domain. In this study, we introduce the novel concept of *consequence-based explanations*, a type of explanation that emphasizes the individual impact of consuming a recommended item on the user, which makes the effect of following recommendations clearer. We conducted an online user study to examine our assumption about the appreciation of consequence-based explanations and their impacts on different explanation aims in recommender systems. Our findings highlight the importance of consequence-based explanations, which were well-received by users and effectively improved user satisfaction in recommender systems. These results provide valuable insights for designing engaging explanations that can enhance the overall user experience in decision-making.

Keywords

recommender systems, consequence-based explanations, decision-making, human-centered computing

1. Introduction

Recommender systems (RS) assist users in decision-making by predicting which items in a catalog will be of interest to them [1]. They help users to make efficient and satisfying decisions with many available options in various domains, including, movie selection [2, 3], meal choices [4, 5], or apartment considerations [6, 7]. The presentation of recommendations is crucial to the overall user experience, and explaining them is an important aspect [8]. Although current explanation approaches have proven useful to extend RS, the need for improvements remains, specifically in terms of generating natural language explanations and explaining impacts [9].


To address these needs, we introduce *consequence-based explanations*, which focus on the impact of consuming recommended items as the main argument. These explanations are particularly beneficial in domains where users have limited knowledge and struggle to understand

IntRS'23: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, September 18, 2023, Singapore (hybrid event)

✉ slubos@ist.tugraz.at (S. Lubos); ttrang@ist.tugraz.at (T. N. T. Tran); sedapolat@gmail.com (S. P. Erdeniz); merfat.el-mansi@student.tugraz.at (M. E. Mansi); afelfernig@ist.tugraz.at (A. Felfernig); manfred.wundara@villach.at (M. Wundara); gerhard.leitner@aau.at (G. Leitner)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

the effects of following recommendations. By emphasizing the actual changes and outcomes resulting from a recommendation, users can gain a clearer understanding of how their choices directly influence their circumstances. To validate this concept, we conducted an online user study to assess users' appreciation of consequence-based explanations and compared different ways of formulating them in two domains with different levels of involvement [10]. Our study provides initial guidance for generating consequence-based explanations.

2. Related Work

Explaining recommendations and providing users with an understanding of why a particular item is recommended emerged as an important topic in recent years [9]. Providing such explanations is critical for shaping the user experience of an RS [8]. Consequently, seven explanation aims have been suggested [11]: (i) *Effectiveness* to help users make good decisions, (ii) *Efficiency* to help users make decisions faster, (iii) *Persuasiveness* to convince users, (iv) *Satisfaction* to improve the user experience, (v) *Scrutability* to allow users to correct the system, (vi) *Transparency* to explain how the system works, and (vii) *Trust* to increase users' confidence in the system. Studies investigating the effect of different methods for generating and visualizing explanations show that there is no optimal approach that fits all dimensions, and explanations should be tailored to the specific goal of the RS [12].

Explanations in RS influence users by persuading them to consume recommended items [13], increasing the perceived usefulness of the RS [14], and contributing to overall user satisfaction [15]. Domain-specific content data enhances explanation effectiveness, while better transparency leads to higher user satisfaction [12]. Users prefer familiar explanation types, as those reduce cognitive effort. However, pure optimization for efficiency is not always useful, as users are willing to spend time analyzing explanations to make good decisions [12].

Explanation types can be categorized depending on how they are generated [9]. *Model-intrinsic* approaches use interpretable recommendation models that directly provide explanations, while *model-agnostic* approaches generate explanations from recommended items. One common technique at this level is the *content-based* style explanation, which utilizes item features and past consumption data for generation [16].

When users have difficulties evaluating recommended options due to limited domain knowledge, mentioning the impact of potential choices explicitly in a *consequence-based* explanation can help. While this explanation type is new to RS, it has proven to be useful in other applications. Reinforcement learning agents explained their behavior in terms of expected consequences of state transitions, to increase human understanding and to enable evaluation of the plausibility of the agents' decisions [17]. Including potential consequences in tornado warnings increased the likelihood of persons taking protective actions [18], and visualizing personal consequences of decisions supported financial planning [19].

3. Consequence-based Explanations

In this paper, we introduce *Consequence-based* explanations, which highlight the positive and negative impacts of consuming a recommended item. This type of explanation emphasizes

the individual effects of choice rather than the underlying factors that led to it. They support users in decision-making, by highlighting how the decision toward an option influences the circumstances in a desired or undesired way. For example, a movie recommender might generate a consequence-based explanation like: *"Toy Story has been recommended to you because it will entertain your whole family, and teach your children about the value of friendship"*. Those explanations can be created using a model-agnostic approach, using the feature descriptors of the recommended items, e.g., the genre and keyword tags. To simplify the decision-making process, it is advisable to include only important consequences for users, considering the multitude of potential outcomes. This aligns with how people typically communicate explanations, focusing on relevant causes rather than overwhelming with excessive information [20].

This paper explores explanations in two domains representing different *levels of involvement* [10]. While we examine *apartments* as high-involvement items, requiring significant time and effort for decision-making due to their long-term and financial impact, we consider *recipes* as a low-involvement domain with shorter-term and less serious impact¹.

Consequences are derived using a rule-based approach customized to each domain, leveraging item features and user preferences (see Section 4.1.1). A sentence is prepared to explain the consequence of each feature value associated with a recommended item, taking into account the specified user preferences. These individual sentences are then combined to form the complete explanation for the recommendation. For instance, the recipe recommender considers user data such as *activity level* and *weight aim* to suggest a recipe with suitable nutrient quantities. By integrating these preferences with nutritional data of recipes, the appropriate sentence for the overall explanation is determined.

We distinguish two types of consequence-based explanations based on their formulation: (i) *motivating consequence-based* explanations formulate the impact in a positive way, expressing which favorable consequence the suggested item has (ii) *avoiding consequence-based* explanations highlight the negative impact that can be prevented by choosing the suggested item. An example for both types is presented in Table 1. To ensure transparency, the downsides of a recommendation are also communicated, if the suggested item does not fulfill all user preferences. For instance, in the apartment domain, an explanation might state that *"your children will need to share bedrooms in this apartment"*.

4. Research Questions and User Study

This study evaluates two formulations of consequence-based explanations and their impact on explanation aims in RS [11]. It compares them to standard content-based explanations as the baseline across two domains with varying levels of involvement. The results offer initial guidance for generating consequence-based explanations that emphasize decision impact. The study is designed and conducted to answer four research questions:

RQ1: How do consequence-based explanations contribute to the explanation aims in RS?

RQ2: How do different formulations of consequences influence their contribution towards the explanation aims?

¹Unhealthy nutrition can have significant consequences, but the impact of a single decision to cook a recipe is low-involvement.

Table 1

Examples of *motivating* and *avoiding* consequence-based explanations, using predefined sentences. The motivating formulation emphasizes the positive impact of consuming the recommended recipe on the users’ activity level and weight aim, while the avoiding formulation highlights the negative impacts on those properties that are prevented with the recommended recipe.

| Motivating Consequence-based | Avoiding Consequence-based |
|--|--|
| The number of carbs, sugar, and protein in the cooked meal will give you enough energy for your activity level, and the number of calories and fat in the dish will support you in losing weight . | The number of carbs, sugar, and protein in the cooked meal will not fall below the needed energy for your activity level, and the number of calories and fat in the dish will not interfere with your aim of losing weight . |

RQ3: How does the level of involvement of items affect the contribution of consequence-based explanations?

RQ4: Is there a correlation between users’ demographics and preferences and the effect on the explanation aims?

Answers to those questions provide insights into (i) user appreciation, (ii) preferences, and (iii) the relevance of consequence-based explanations across domains. This helps improve user support in RS with more valuable explanations.

4.1. User Study

4.1.1. Dataset Preparation

To prepare the user study, we created small datasets consisting of 20 items per domain. These datasets were designed with feature descriptors that included the necessary properties for generating explanations. For the recipe domain, data from recipe websites² was collected and processed in a uniform format including *title*, *description with ingredients*, and a *generated image*³ using the recipe title as input. Features include the *cuisine type*, *difficulty*, *diet*, *cooking time*, *nutritional data*, and *allergenic*. For the apartment dataset, *size*, *rent*, *number of bedrooms*, *distance to the city center*, *availability of private parking and garden*, and *distance to leisure facilities* are used as features. Those samples have been created such that the content feature values are equally distributed. Those have been extended with a *generated photo*, *title*, and *description* in the style of an apartment advertisement⁴.

4.1.2. Generation of Recommendations

Since this paper focuses on explanations, the implemented RS is simple but generates personalized recommendations based on user preferences and item features. Preferences in the recipe domain, include *favorite cuisine*, *followed diets*, *preferred cooking time*, *cooking skill*, *ingredients to be avoided*, *activity level*, and *weight aim*. In the apartment domain, preferences include the

²<https://www.aheadofthyme.com>, <https://www.seriousseats.com>, <https://www.themediterraneanandish.com>, and <https://theplantbasedschool.com>

³The photo was created with DALL-E 2, an AI image generator (<https://openai.com/product/dall-e-2>)

⁴Descriptions were generated with ChatGPT using features as input. Their correctness has been inspected manually.

number of children living in the apartment, rent, preferred distance to the city center, availability of a private car, and favorite leisure time activities. Recommendations are generated in a two-step process of candidate generation and ranking [21]. Firstly, the dataset is filtered for candidates fulfilling strict user preferences, i.e., diet, ingredients, cooking time, and cooking skill for the recipe domain, and city center distance and rent for the apartment domain. The remaining candidates are then ranked by compatibility with the *Multi-Attribute Utility Method (MAUT)* [22], assigning scores to item alternatives, depending on their overlap with feature preferences.

4.1.3. Study Design

We conducted an online study adopting the design used in [8]. Participants were randomly assigned to either of the domains, where they were presented with a scenario of deciding on a recipe to cook later that day or searching for an apartment to move in for the next years, based on their personal situation and interests. After providing demographic information (age, gender, education) and preferences (see Section 4.1.2), a suggested item was presented, without any content information, only described by the generated explanation. Participants were evenly distributed among the different explanation variants (motivating or avoiding consequence-based and content-based). They used a *5-point Likert scale* to rate their likelihood of choosing the suggested item, along with their perception of the explanation’s satisfaction and understandability. Furthermore, the importance of individual consequences was rated to evaluate suggestion quality. In the final step, participants received another suggestion with a complete item description and rated their likelihood of considering it. Notably, participants were unaware that the same item was suggested twice with different descriptions, enabling the assessment of explanation effects by comparing rating differences.

5. Evaluation

5.1. Study Participants

In total, 114 persons participated in the study. 63 were part of the recipe domain, while 51 used the apartment recommendation. Around 25% of participants identified as female. In general, the participants were rather young, with an average of 26 years. Influenced by this, 51% of the participants specified *high school* as their highest educational degree, while 45% had graduate degrees. These demographic factors need to be considered to interpret the results.

5.2. Analysis Methods

After gathering participant responses, we evaluated the impact of consequence-based explanations on the explanation aims [11] using metrics proposed in [12], comparing the results to a content-based explanation baseline. The collected datasets share common features, such as non-normal distribution⁵, contain ordinal values, and violate independence due to the within-subjects study design. Therefore, non-parametric tests were used in *SPSS* as hypothesis tests with $\alpha = 0.05$. In case the hypothesis test showed significance, the value of α was updated for

⁵following the *Shapiro-Wilk* test, with $\alpha < 0.05$

pairwise follow-up tests as $\alpha_{new} = \frac{\alpha}{N}$, where N is the number of pairwise follow-up tests of the respective attributes, to avoid error type 1.

5.3. Explanation Aims

The user study aims to investigate the impact of consequence-based explanations on explanation aims in RS [11], comparing different formulations in domains with varying levels of involvement. It assesses efficiency, effectiveness, persuasiveness, satisfaction, and transparency, but does not examine the impact on trust and scrutability.

5.3.1. Efficiency

The study analyzed the *item-based efficiency*, i.e., the time users spend to evaluate an item and decide on the likelihood of consuming it [12], of consequence-based explanations compared to content-based explanations in low and high-involvement domains, using recipes and apartments as representatives. The results (see Table 2a) indicate that consequence-based explanations were more efficient, as users took less time to decide, especially with motivating consequences. The *Kruskal-Wallis* test and follow-up *Mann-Whitney-U* tests confirmed that those were significantly more efficient than avoiding consequences and the baseline in the recipe domain. However, there was no significant difference between avoiding consequences and content-based explanations or between explanation variants in the apartment domain. The better efficiency of consequence-based explanations suggests that including individual impact in the explanation helps users decide faster. The slight advantage of motivating consequences could be due to users' interest in the positive effects of a decision or lower cognitive effort if negations are avoided.

Comparing the domains, the item-based efficiency was slightly better in the apartment domain, without statistical significance. The duration differences between the motivating and avoiding consequence formulation were small within the domains. Avoiding consequences were more efficient in the apartment domain, likely because pointing out the serious avoided negative consequences helps users make faster decisions. In contrast, in the recipe domain, motivating consequences were more efficient, probably because the positive impact of consuming food is more relevant in this case.

Overall, the analysis shows that consequence-based explanations improve decision-making efficiency. Considering, that they were longer (338-749 characters) compared to the content-based explanations (186-524 characters), the effect on effectiveness is further strengthened. The motivation of positive consequences seems to be more efficient in low-involvement domains, while the avoidance of negative impact tends to be more efficient in high-involvement domains. This aligns with the *prospect theory* [23], indicating that people are risk-averse with gains but risk-seeking with losses. People tend to avoid loss in important decisions.

5.3.2. Effectiveness and Persuasiveness

To measure effectiveness, the ratings for the explanation-based and content-description-based recommendations of the same user are compared [8, 24]. Similar ratings indicate effectiveness, as both ways of presenting the suggestion led to the same result. If the full content description of a recommendation receives lower ratings than the explanation, users tend to overestimate

Table 2

Mean study results by explanation aims [11]. The efficiency results represent the average time it took users to decide the likelihood of consuming an item based on the explanation. Effectiveness is calculated as the difference between the ratings given by a user for the explanation and the content description. A value close to zero indicates effectiveness, while a positive value suggests persuasiveness. For satisfaction and transparency, mean values of user-perceived ratings are shown.

| Domains | Motivating | | | Avoiding | | | Content-Based | | | Overall | | |
|----------------|------------|-------|-------|----------|-------|-------|---------------|-------|-------|---------|-------|-------|
| | Apt | Rec | All | Apt | Rec | All | Apt | Rec | All | Apt | Rec | All |
| Efficiency [s] | 132.9 | 130.5 | 131.6 | 125.7 | 138.8 | 133.0 | 173.1 | 147.5 | 159.3 | 143.9 | 138.8 | 141.1 |
| Effectiveness | 0.4 | 0.1 | 0.2 | 0.2 | -0.3 | -0.1 | 0.2 | -0.6 | -0.3 | 0.2 | -0.3 | -0.1 |
| Satisfaction | 3.8 | 3.4 | 3.6 | 3.8 | 3.7 | 3.7 | 3.5 | 3.3 | 3.4 | 3.7 | 3.5 | 3.6 |
| Transparency | 3.9 | 3.6 | 3.8 | 3.8 | 3.7 | 3.7 | 3.8 | 4.0 | 4.0 | 3.8 | 3.8 | 3.8 |

(a) **Explanation types and formulations**, grouped by domains.

| Groups | Gender | | Education | | Cooking Time | | | Activity Level | | | | Weight Aim | | |
|----------------|--------|-------|-----------|-------|--------------|---------|-------|----------------|--------|-------|-------|------------|-------|-------|
| | F | M | Non-uni. | Uni. | No Pref. | <30 min | <1h | Sed. | Light. | Mod. | Very | Lose | Mnt. | Gain |
| Efficiency [s] | 147.5 | 133.2 | 126.9 | 146.4 | 165.6 | 132.1 | 124.3 | 143.9 | 126.0 | 128.0 | 183.4 | 131.2 | 151.2 | 121.0 |
| Effectiveness | - | 0 | - | - | - | 0.1 | - | - | - | 0.2 | -1 | - | - | 0.4 |
| | 1.2 | | 0.2 | 0.4 | 0.9 | | 0.3 | 0.4 | 0.4 | | | 0.6 | 0.2 | |
| Satisfaction | 3.6 | 3.4 | 3.4 | 3.6 | 3.7 | 3.5 | 3.4 | 3.7 | 3.5 | 3.2 | 3.7 | 3.7 | 3.3 | 3.7 |
| Transparency | 4.0 | 3.7 | 3.5 | 4.1 | 3.6 | 3.8 | 3.8 | 3.6 | 4.1 | 3.4 | 4.3 | 3.9 | 3.7 | 4.0 |

(b) **Low-involvement** domain (recipes), grouped by demographics and preferences (excluding *preferred diet* due to even distribution).

| Groups | Gender | | Education | | Rent | | | City Center Distance | | | Car Availability | |
|----------------|--------|-------|-----------|-------|-------|-----------|-----------|----------------------|-------|--------|------------------|-------|
| | F | M | Non-uni. | Uni. | <500€ | 500€-700€ | 700€-900€ | <1km | 1-5km | 5-10km | Yes | No |
| Efficiency [s] | 147.3 | 145.1 | 116.8 | 147.3 | 141.0 | 145.7 | 144.2 | 164.2 | 112.1 | 182.1 | 163.5 | 127.8 |
| Effectiveness | 0.5 | 0.2 | 0.6 | - | - | 0.5 | 0.3 | - | 0.5 | 0.1 | 0.3 | 0.2 |
| | | | 0.4 | 0.2 | 0.2 | | | 0.3 | | | | |
| Satisfaction | 3.8 | 3.6 | 3.8 | 3.5 | 4 | 3.5 | 3.6 | 3.9 | 3.7 | 3.5 | 3.6 | 3.8 |
| Transparency | 3.8 | 3.9 | 4.1 | 3.5 | 4.2 | 3.9 | 3.5 | 4.3 | 4.0 | 3.3 | 4.0 | 3.7 |

(c) **High-involvement** domain (apartments), grouped by demographics and preferences (excluding *children* due to even distribution).

the explanation, indicating persuasiveness. On the contrary case, users tend to underestimate the explanation, indicating weak persuasiveness.

The descriptive statistics (see Table 2a) show that all observed explanation types, including the baseline, were effective, with minimal mean rating differences. This is expected as the necessary information to decide is already included in the explanation. However, motivating consequences showed a trend of slight persuasiveness, while avoiding consequences and the baseline tended to be underestimated. Consequence-based explanations were effective for both low- and high-involvement item domains. For apartments, users found motivating consequences slightly more persuasive, resembling advertising tactics that highlight positive aspects. For recipes, users tended to underestimate the item, particularly when avoiding consequences were mentioned, likely because negative consequences are generally weak in this domain.

The analysis suggests that consequence-based explanations provide an effective way to explain recommendations in both low- and high-involvement item domains, where the motivating formulation leads to slight persuasiveness.

5.3.3. Satisfaction

To assess user satisfaction with the explanations, participants were asked to rate their satisfaction with the recommended item [24, 12]. Consequence-based explanations resulted in higher user satisfaction compared to the baseline in both low- and high-involvement domains. Users appreciated the clear communication of the suggested item's impact, with avoidance of consequences leading to the highest user satisfaction. Aligned with the *prospect theory* [23], it indicates that users tend to prioritize avoiding negative impacts when making choices based on the explanation.

Based on the descriptive statistics (see Table 2a), we found that user satisfaction was higher in the high-involvement domain than in the low-involvement domain, which supports our assumption and may be explained due to the higher criticality of the decision's impact. While no difference was observed between motivating and avoiding consequences in the apartment domain, avoiding consequences received higher ratings in the recipe domain. Overall, the results indicate that consequence-based explanations are well-received by users and positively affect user satisfaction with the RS.

5.3.4. Transparency

To evaluate the *user-perceived transparency* of the explanations [24, 12], participants rated how well the explanation helped them understand the reasoning behind the recommendation. The descriptive statistics indicate (see Table 2a) that consequence-based explanations, in both formulations and domains, did not excel in terms of transparency as they lack insights into how the system generated the suggestion.

5.3.5. Summary

In response to RQ1, RQ2, and RQ3, the study found that consequence-based explanations are efficient for explaining recommendations in both domains. Motivating consequences were more efficient in the low-involvement domain, while mentioning avoided negative consequences was more efficient in the high-involvement domain. Both formulations were effective, with motivating consequences being slightly more persuasive than avoiding consequences. Users expressed higher satisfaction compared to content-based explanations, with the avoiding formulation leading to the highest satisfaction. In terms of transparency, consequence-based explanations did not outperform the baseline.

5.4. Influence of Demographics and User Preferences

RQ4 explores the connection between users' demographics, preferences, and their impact on the explanation aims. Descriptive statistics for both domains are presented in Tables 2b and 2c. The analysis excluded age due to participant homogeneity. Persons with higher

education levels spent more time analyzing explanations in both domains, likely due to their tendency for thorough decision-making. In the recipe domain, participants with specific requirements (e.g., preferred cooking time, weight aim) evaluated explanations faster and found them more transparent, showing a better understanding of the recommendation process. Female participants, those without a preference for cooking time, and very active individuals tended to underestimate the recommended recipe. This suggests that more detailed and persuasive explanations could be beneficial for those. For apartments, participants without a university degree found the explanations more persuasive, while graduates tended to underestimate the recommended item. This indicates that graduates may prefer more detailed and comprehensible explanations. Conclusions for user satisfaction are not drawn due to homogeneous results.

5.5. Importance of Features

To determine which features in consequences influence users' decisions to consume a recommended item, participants rated their importance. Analyzing the data using the *Kruskal-Wallis* test showed significant differences across the features in both domains, which have been identified with pairwise *Mann-Whitney-U* tests. The results indicate that consequences with a strong personal impact were valued higher than those describing the item more generally. In the apartment domain, *financial liability* is an important factor to be included, with the *monthly rate* consequence being more important than the *distance to the city center*, a *private garden*, a *private parking spot*, *number of bedrooms*, and *distance to preferred leisure activities*. Furthermore, the *number of bedrooms* and the *distance to the city center* are more important than having a *garden*. In the recipe domain, *preparation time* is more important than the *cuisine type*, *difficulty level* of a recipe, and individual *activity level*. The *weight aim*, *preferred diet*, and *avoidance of ingredients* are more important than the *activity level* consequence.

6. Conclusions, Limitations, and Future Work

This paper introduced and evaluated consequence-based explanations as a new explanation type for recommendations, emphasizing the individual impact of a recommended item. Our user study confirmed user appreciation, showing a trend toward increased satisfaction. These explanations support users in making effective and efficient decisions, particularly in high-involvement domains where the impact is crucial, and users may lack expertise. By highlighting personal impact, valuable insights for understanding the reasons behind recommendations and making informed decisions are provided.

The initial study analyzes this novel explanation type, while it acknowledges certain limitations that will be addressed in future work. One limitation is the lack of statistically significant results for some measured dimensions. A reason might be the small number of participants in the study. To overcome this, we plan to conduct a larger study with a more diverse group of participants, in terms of age, gender, and educational background, to gain a deeper understanding of the findings. Another limitation is the use of synthetic datasets, which may impact the evaluation of recommended items. To address this concern, we intend to validate the consequence-based explanations in a real-world study, which will provide practical insights. Finally, the current methods for generating our explanations are limited to the recommended

item content. In the future, we aim to enhance this method by incorporating collaborative data and refining personalization techniques. Additionally, we expect the integration of generative AI to offer diverse and natural explanations, ensuring their authenticity and eliminating any false consequences, as a beneficial possibility for improving the explanations.

Acknowledgments

The presented work has been developed within the research project STREAMDIVER which is funded by the Austrian Research Promotion Agency (FFG) under the project number 886205.

References

- [1] F. Ricci, L. Rokach, B. Shapira, *Recommender Systems: Techniques, Applications, and Challenges*, Springer US, New York, NY, 2022, pp. 1–35. doi:10.1007/978-1-0716-2197-4_1.
- [2] C. A. Gomez-Uribe, N. Hunt, The netflix recommender system: Algorithms, business value, and innovation, *ACM Trans. Manage. Inf. Syst.* 6 (2016). doi:10.1145/2843948.
- [3] N. Sharma, M. Dutta, Movie recommendation systems: A brief overview, in: *Proceedings of the 8th International Conference on Computer and Communications Management, ICCCM '20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 59–62. doi:10.1145/3411174.3411194.
- [4] M. Ge, F. Ricci, D. Massimo, Health-aware food recommender system, in: *Proceedings of the 9th ACM Conference on Recommender Systems, RecSys '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 333–334. doi:10.1145/2792838.2796554.
- [5] F. Pecune, L. Callebert, S. Marsella, *A recommender system for healthy and personalized recipe recommendations*, 2020.
- [6] R. Burke, Knowledge-based recommender systems, *Encyclopedia of library and information systems* 69 (2000) 175–186.
- [7] A. Gharahighehi, K. Pliakos, C. Vens, Recommender systems in the real estate market—a survey, *Applied Sciences* 11 (2021). doi:10.3390/app11167502.
- [8] M. Bilgic, R. J. Mooney, Explaining recommendations: Satisfaction vs. promotion, in: *Beyond personalization workshop, IUI*, volume 5, 2005, p. 153.
- [9] Y. Zhang, X. Chen, Explainable recommendation: A survey and new perspectives, *Found. Trends Inf. Retr.* 14 (2020) 1–101. doi:10.1561/1500000066.
- [10] A. Felfernig, M. Atas, T. N. T. Tran, M. Stettinger, S. P. Erdeniz, G. Leitner, An analysis of group recommendation heuristics for high-and low-involvement items, in: *Advances in Artificial Intelligence: From Theory to Practice: 30th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2017*, Arras, France, June 27-30, 2017, *Proceedings, Part I* 30, Springer, 2017, pp. 335–344.
- [11] N. Tintarev, J. Masthoff, A survey of explanations in recommender systems, in: *2007 IEEE 23rd international conference on data engineering workshop, IEEE*, 2007, pp. 801–810.
- [12] F. Gedikli, D. Jannach, M. Ge, How should i explain? a comparison of different explanation types for recommender systems, *International Journal of Human-Computer Studies* 72 (2014) 367–382.

- [13] S. Gkika, G. Lekakos, The persuasive role of explanations in recommender systems., in: BCSS@ PERSUASIVE, 2014, pp. 59–68.
- [14] M. Zanker, D. Ninaus, Knowledgeable explanations for recommender systems, in: 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, volume 1, 2010, pp. 657–660. doi:10.1109/WI-IAT.2010.131.
- [15] T. N. T. Tran, V. M. Le, M. Atas, A. Felfernig, M. Stettinger, A. Popescu, Do users appreciate explanations of recommendations? an analysis in the movie domain, in: Fifteenth ACM Conference on Recommender Systems, 2021, pp. 645–650.
- [16] N. Tintarev, J. Masthoff, Beyond Explaining Single Item Recommendations, Springer US, New York, NY, 2022, pp. 711–756. doi:10.1007/978-1-0716-2197-4_19.
- [17] J. van der Waa, J. van Diggelen, K. van den Bosch, M. Neerincx, Contrastive Explanations for Reinforcement Learning in terms of Expected Consequences, arXiv e-prints (2018). doi:10.48550/arXiv.1807.08706.
- [18] J. T. Ripberger, C. L. Silva, H. C. Jenkins-Smith, M. James, The influence of consequence-based messages on public responses to tornado warnings, *Bulletin of the American Meteorological Society* 96 (2015) 577–590.
- [19] A. E. Fano, S. W. Kurth, Personal choice point: helping users visualize what it means to buy a bmw, in: International Conference on Intelligent User Interfaces, 2003.
- [20] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, *Artificial Intelligence* 267 (2019) 1–38. doi:10.1016/j.artint.2018.07.007.
- [21] J. Davidson, B. Livingston, D. Sampath, B. Liebald, J. Liu, P. Nandy, T. Van Vleet, U. Gargi, S. Gupta, Y. He, M. Lambert, The YouTube video recommendation system, in: Proceedings of the fourth ACM conference on Recommender systems - RecSys '10, ACM Press, Barcelona, Spain, 2010, p. 293. doi:10.1145/1864708.1864770.
- [22] S. Jansen, The Multi-attribute Utility Method, 2011, pp. 101–125. doi:10.1007/978-90-481-8894-9_5.
- [23] D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk, *Econometrica* 47 (1979) 263–291.
- [24] N. Tintarev, J. Masthoff, Evaluating the effectiveness of explanations for recommender systems: Methodological issues and empirical studies on the impact of personalization, *User Modeling and User-Adapted Interaction* 22 (2012) 399–439.