The Effect of Relational versus Anecdotal Explanations in Movie Domain Recommendations

Liam de la Cour, Derek Bridge

School of Computer Science & Information Technology, University College Cork, Ireland

Abstract

This paper explores the effect *explanation type* has on people's perceptions of recommendation quality. More specifically, we explore this effect in systems whose users 'consume' an entity by reading about the entity. In such systems, one of the main goals is to persuade the user to extend their exploration of the domain. We compare two explanation types: *relational explanations* and *anecdotal explanations*. We compare them in the movie domain using a between-subject study. We use Path Analysis (PA) to evaluate our results. We find that using anecdotal explanations positively affects how *informative* and *entertaining* participants find explanations, which, in turn, positively impacts how *interesting* the user finds the explanation. Finally, this positively affects the perceived quality of the recommendations. We also explore the impact a user's level of *domain engagement* has on these factors. We find that it positively correlates with how interesting they perceive the explanations to be and with the perceived recommendation quality.

Keywords

recommender systems, explanations, path analysis

1. Introduction

Recommender Systems (RSs) suggest items for users to consume, based on inferred tastes, e.g. movies to watch, songs to listen to, paintings to look at, etc. In this work, however, we are working in systems that allow users to become informed about a domain. In these systems, we refer to the items as Information Entities (IEs). In some cases, an IE will correspond to items the user can consume in the real world (movies, songs, works of art, etc.). But in other cases, they will correspond to real-world things that users would not normally consume (the director of a movie, the recording artist of a song); they might even refer to abstract concepts (the period in which a work of art was produced, or the artistic movement to which it belongs). Crucially, each IE has some information associated with it – for now, we assume it is a piece of text, such as a biography if the IE corresponds to a person, a synopsis if the IE corresponds to a movie, and so on. Consumption in these systems is done by reading the text, to learn about the IE. Of course, it may be that, after consuming the IE (i.e. reading about it), the user does consume the corresponding real-world item (e.g. watch the corresponding movie), where this is possible. Equally, since IEs in these systems are highly inter-connected, the user may wish to continue exploring the domain by reading information about other, related IEs.

The users of these systems are motivated to explore the domain. They may be less casual than the users of a conventional RS. They may even be enthusiasts, who already know the domain to some degree, but wish to learn more. However, the number of IEs will be large — arguably, larger than the number of items in a conventional RS since there will be IEs that correspond to conventional items but IEs that correspond to other entities and concepts too. These system need a RS to assist exploration of the domain. In this paper, our focus is on the explanations that we might use in this kind of RS.

We review explanations in more depth in Section 2. But, in overview, the kind of explanation we look at in this paper takes the form of a short text (typically, one or two sentences), that connects IEs, e.g., from the user's profile, to a recommended IE. The goal of the explanation is not one of transparency; rather, it is primarily one of persuasion — to encourage the user to consume (read about) the recommended IE.

IntRS'24: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, October 18, 2024, Bari (Italy)
liamdelacour@umail.ucc.ie (L. de la Cour); d.bridge@cs.ucc.ie (D. Bridge)
0009-0008-1043-6277 (L. de la Cour); 0000-0002-8720-3876 (D. Bridge)

Persuasive explanations can help the motivated user explore and learn more about the domain.

More specifically, this paper reports an experiment we have conducted to investigate the effect explanations have on perceived recommendation quality. We wish to find whether certain types of explanations can increase perceived recommendation quality, and what aspects of these explanations drive this improvement. We evaluate this with a user study. To conduct the study, we have built a system called MovieBuff. It allows movie enthusiasts to learn more about movies, crew members, filming locations, studios, etc. In the experiment, we limit ourselves to recommendations about movies and crew (i.e. people such as actors, directors, and so on). As we will explain in detail in Section 3, participants receive recommendations for movies and for crew members. Explanations connect the crew member recommendations to the movie recommendations. The explanations are of two types — *relational* and *anecdotal* — and our goal is to compare the effects of these two types of explanations.

Our research questions are as follows:

- **RQ1**: Which of these explanation types (*relational* or *anecdotal*) increases perceived recommendation quality?
- **RQ2**: How does the informativeness, entertainment value, and interestingness of the explanations affect perceievd recommendation quality?
- **RQ3**: What impact does an individual user's domain engagement and item familiarity have on perceived recommendation quality?

We organise this work as follows: Section 2 briefly surveys current work in this area and how we fit into the landscape; Section 3 presents the MovieBuff user study; Section 4 discusses the results; finally, in Section 5, we conclude and discuss future work.

2. Related Work

A number of papers describe systems of the kind we are interested in — ones where consumption means reading about an entity. For example, Durao & Bridge describe a browser that allows its users to explore a linked data graph [1]. Nodes in the graph are highly inter-connected. So the browser incorporates a RS, which uses a novel classifier to predict whether a node will be of interest to the user or not, based on the user's previous browsing behaviour. In [2], the emphasis is on a form of story-telling, rather than browsing. Users supply start and end IEs and De Vocht *et al.* use an A* algorithm to find a path that connects the corresponding nodes. Path quality is enhanced by filtering edges for relevance, and by defining edge weights and node heuristics. A third, different approach is to support exploration by summarizing facts about IEs. In particular, Nuzzolese *et al.* mine patterns from linked data, where a pattern captures the most relevant facts for describing an entity of a given type [3]. Given a specific entity, instantiating the pattern for that type of entity is a way of selecting, organizing and visualizing knowledge about that specific entity.

In ths paper, we are interested in the role of explanations in this kind of system. Much work has been done in recent years to explore the effects explanations have on users of Recommender Systems in general. Tintarev & Masthoff [4] list seven goals that explanations may have, namely: Transparency, Scrutability, Trust, Effectiveness, Persuasiveness, Efficiency, and Satisfaction. Much work has shown that the type of explanations used in a RS can influence these goals [5, 6, 7, 8, 9, 10, 11, 12]. However, explanation types that lead to an improvement in one goal may have a negative effect on another goal. For example, Tintarev & Masthoff [6] found that personalising explanations was detrimental to effectiveness but positively affected satisfaction. Bilgic & Mooney [8] found that certain styles of explanation caused users to over-estimate the relevance of the recommended item. They found that these styles of explanation were persuasive but led to lower satisfaction. In our work, we will explore how our particular explanation types affect persuasion and satisfaction.

Explanations in RSs have been explored in domains such as music [13, 14], movies [15, 6, 16, 9, 17, 18, 19, 20, 12], books [8], cultural heritage [11, 21], academia [22, 10, 7, 23], tourism [24, 25] and e-commerce [5, 26]. Our work is undertaken in the movie domain. However, as explained already, we do not recommend movies for the user to watch. Instead, we recommend movies and their crew members, along with information about these, so that motivated users can learn more about the domain.

Explanations can be categorised in many ways. Friedrich & Zanker [27] categorize them on three dimensions: the type of information exploited to create the explanation (item information, user information or other information); the paradigm (collaborative, content-based or knowledge-based); and the reasoning model (white box or black box). Radensky *et al.* [22] categorise by scope: global explanations explain a model's overall decision-making process; local explanations explain individual recommendations. Explanations can also be classified by how they are presented to the user. For example, they may be text-based or they may use visualizations. Zanker & Schoberegger [5] explore how various text-based explanations can affect persuasiveness. They include fact-based, argumentative fact-based and argumentative sentences in their study. Hernandez-Bocanegra & Ziegler [24] explore text-based explanations in conversational agents to see what kinds of explanations users want to receive in these systems. Herlocker *et al.* [16] explore twenty explanation methods, and found that some of their histogram explanations had the most positive effect on users. Gedikli *et al.* [9] use tag clouds for explanation and find that these increase both transparency and satisfaction for the user.

In terms of the above categorization, MovieBuff's explanations: use item information; are content-based; are white-box; are local; and are text-based.

Since explanations affect users' perceptions, explanation types and their effects must be evaluated through user studies. In the literature, user studies gather data about explanations through the use of questionnaires, where participants provide, e.g., ratings [16], re-ratings [6], or rankings [5]; other studies use interviews [21]. In some work, background information is also collected, such as the user's level of domain expertise [23], since different explanations may be preferred by users of varying expertise. Behavioural metrics, such as time spent using the system [7], can also be gathered and are useful because they do not require the participant to explicitly answer questions about the explanations.

Work in this area uses a variety of statistical methods to evaluate hypotheses. These techniques include two-tailed t-tests [5], ANOVA analysis [16], Path Analysis (PA) [9], and Structural Equation Modeling (SEM) [7, 14, 23]. In this paper, we will follow the Knijnenburg *et al.* Evaluation Framework [28] and use PA to evaluate our results, as it allows us to determine which aspects of the explanations cause the observed effects.

3. User Study

To answer our research questions, we conducted a between-subject user study in the movie domain. We start by describing the data we use.

3.1. Movie Data

We scraped movie data from IMDb.[1] The API endpoint we used to scrape this data allowed us to return the top 100 from IMDB's 'Most Popular' section, as well as an equivalent for each of 12 genres. Scraping this resulted in 917 distinct movies. For each of the 917 movies, we obtained the title, a synopsis and an image; and for each of the actors & other crew members, we obtained their name, biography and, where possible, an image. We also scraped IMDb's movie trivia. Each piece of movie trivia is, loosely speaking, an anecdote about the movie. We use the trivia in one of our explanation types (see Section 3.3). Movies and crew members are our IEs; movie synopses and crew member biographies are what users consume when they use MovieBuff to explore the domain. Details are in Table 1.

We store the movie data as a knowledge graph. Each movie, crew member and anecdote is a node in the graph. Labeled edges connect movies and their crew members. Edges also connect each anecdote with the movie from whose IMDb page it was scraped. Additionally, we apply Named Entity Recognition to the text of the anecdote to find names and, where possible, we connect anecdote nodes to the crew members mentioned in the anecdote. We tokenize the movie synopses, the crew member biographies

[1] https://www.imdb.com/

Table 1
Scraped Movie Data

Node Type	Count
Movies	917
Crew Members	5,687
Movie Trivia	78,540

and the anecdotes; we discard stopwords; and then we apply TF-IDF vectorization. Thus, each of these texts is now represented by a vector of TF-IDF scores.

3.2. Study Procedure

We invited people to participate, telling them that it was a study about movie recommendations & explanations, and that they would be interacting with a system called MovieBuff.

After consenting to participate in the study, participants are assigned at random to one of two groups, and then each participant completes three stages:

Stage one: MovieBuff administers a short questionnaire to measure the participant's level of *domain engagement*. The questions we ask are based on ones from the Goldsmiths Musical Sophistication Index (Gold-MSI) [29]. Gold-MSI comprises five aspects. Four of the aspects are concerned with music skills, such as musical training and singing abilities. We restrict ourselves to the remaining aspect, the one called Active Engagement (AE). AE covers "a range of active musical engagement behaviours (e.g. I often read or search the internet for things related to music) as well as the deliberate allocation of time and money on musical activities (e.g. I listen attentively to music for n hours per day)" [29]. We modify the AE questions so that they refer to movies, instead of music. There are ten questions, each with answers on a seven point scale. We measure a user's domain engagement as the mean of their responses to the ten questions.

Stage two: MovieBuff then requires the participant to browse or search its movie catalog in order to find ten movies that they have previously enjoyed. MovieBuff adds them to a *user profile*.

Stage three: Next, there are five rounds of recommendations. In each round, MovieBuff makes a personalised movie recommendation, displaying the title, an image and the movie synopsis. It also recommends a crew member associated with that movie. It displays the person's name, their biography and their image, if there is one. It also displays an explanation of the relevance of the crew member that it is recommending. One group of participants always sees relational explanations; the other group always sees anecdotal explanations. MovieBuff requires participants to answer questions about the recommended movie, the recommended crew member and the explanation that was shown. An example of what a participant sees is given in Figure 1.

We give more details about Stage three in the next three sections: the explanation types, the way recommendation works, and the questions that the participants must answer.

3.3. Explanation Types in MovieBuff

This work compares two types of explanations, namely *relational* explanations and *anecdotal* explanations. Examples of these can be seen in Table 2.

A relational explanation is simply a statement of the role a crew member played in a movie (e.g. actor, director, screenwriter, etc.). An anecdotal explanation is one of the pieces of trivia that we scraped from IMDb (Section 3.1).

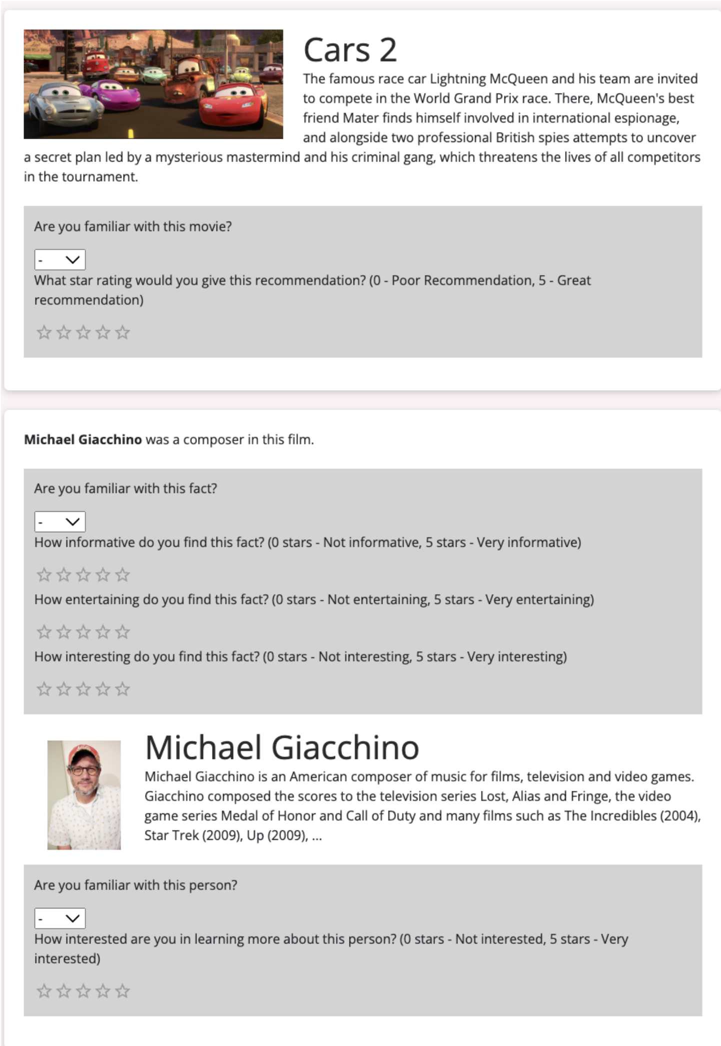

Figure 1: Screenshot of stage three in MovieBuff. In this example, there is a relational explanation ("Michael Giacchino was a composer in this film").

Table 2
Examples of relational explanations and anecdotal explanations

Relational Explanation	Anecdotal Explanation
Dwayne Johnson was a producer of this film	**Dwayne Johnson** and Kevin Hart credit their onscreen chemistry to their offscreen friendship
Tim Allen was a voice actor in this film	The toolbox on top of the milk crate that Woody is trapped in is a Binford, the same tool brand that **Tim Allen** used on his television show Home Improvement (1991).

3.4. Recommendation & Explanation in MovieBuff

In our user study, we want to examine the effect that different types of explanation have on perceived recommendation quality. Recommendation quality itself is not our primary focus — only whether it is affected by explanation type. Hence, we use a very simple content-based recommender.

In each round, MovieBuff recommends a movie. The candidates are movies that are not in the user profile. MovieBuff computes the cosine similarity between each candidate and each movie in the user's

Table 3
Questions asked to participants in each round of recommendations (PC = Personal Characteristic; EXP = User Experience; SSA = Subjective System Aspect. See Section 4.2)

Question	Response Type	Variable Name
Are you familiar with this movie? (**PC**)	Yes/No	Movie Familiarity
What star rating would you give this recommendation? (**EXP**)	5-point scale	Movie Recommendation Quality
Are you familiar with this fact? (**PC**)	Yes/No	Explanation Familiarity
How informative do you find this fact? (**SSA**)	5-point scale	Perceived Informativeness
How entertaining do you find this fact? (**SSA**)	5-point scale	Perceived Entertainment
How interesting do you find this fact? (**SSA**)	5-point scale	Perceived Interestingness
Are you familiar with this person? (**PC**)	Yes/No	Crew Member Familiarity
How interested are you in learning more about this person? (**SSA**)	5-point scale	Crew Member Recommendation Quality

profile. The highest of these similarities determines the movie to be recommended.

MovieBuff also recommends a crew member. The candidates are crew members associated with the recommended movie. If MovieBuff were a deployed RS, then we might select from among the candidates again using cosine similarity either with movies in the user profile or their associated crew members. But, in this user study, the crew member recommendation must come with an explanation — for some users the explanation will be relational and for others it will be anecdotal. It is the effect of this explanation that we are studying. Hence, for the purposes of running the user study, MovieBuff operates in a counter-intuitive way: it recommends the crew member who has the best anecdotal explanation. It works in this way irrespective of whether the user is one of those who will ultimately be shown the anecdotal explanation or will ultimately be shown the relational explanation.

It remains to be said what we mean by the best anecdotal explanation. MovieBuff chooses an anecdote that is connected to the recommended movie whose cosine similarity with the user's profile is highest.

3.5. Question for Study Participants

Participants are asked questions during each round of recommendations to gauge their opinions. These questions can be seen in Figure 1 and are also listed in Table 3 with some additional information. In the case of the movie recommendation, we ask users what star-rating they think they would award this movie on the basis of its synopsis (*recommendation quality*). We ask the same question about the recommended crew member. In the case of the explanation, we ask participants how *informative*, *entertaining* and *interesting* they find it. We also ask whether they are already familiar with the recommended movie, recommended crew member and the content of the explanation.

4. Results

The online experiment ran in January/February 2024 with 106 participants.

Table 4
Results Summary (n=106). Yes/No responses encoded as Yes = 1, No = 0

Variable	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 75	Max
Domain Engagement	0.56	0.14	0.14	0.48	0.63	0.94
Movie Familiarity	0.84	0.2	0.2	0.8	1	1
Movie Recommendation Quality	3.5	0.84	1.4	3	4.2	5
Explanation Familiarity	0.23	0.23	0	0	0.4	0.8
Perceived Informativeness	3	1.1	1	2.2	3.8	5
Perceived Entertainment	2.5	0.94	1	1.8	3.2	5
Perceived Interestingness	2.5	0.95	1	1.8	3.2	4.6
Crew Member Familiarity	0.49	0.24	0	0.4	0.6	1
Crew Member Recommendation Quality	2.4	0.83	1	1.8	3	4.2

4.1. Participants

The participants in this study were a convenience sample of students from our university, recruited through college mailing lists. In the recruiting email, they were informed about the topic of the study. Participants had to confirm on the platform that they were students over 18. The Social Research Ethics Committee (SREC) in our university granted ethical approval for this work (application number 2023-115). To get approval to use MovieBuff to gather results, we were put under the restriction that we could gather no personal information about the participants. As such, we cannot provide a breakdown of participant characteristics, such as age or gender.

Participants were randomly assigned to the group who received relational explanations or the group who received anecdotal explanations. They could withdraw from participating at any time. If they did so, no results from them were recorded. By the end of the experiment, we had results for 45 participants in the relational explanations group, and 61 in the anecdotal explanations group. The imbalance may be due to chance or to more people withdrawing from the less engaging version of the system.

Each participant was presented with five rounds of recommendations, as described in Section 3. All responses were included in the analysis. To avoid the issue of *correlated errors* [30], we average each participant's responses to create a dataset of 106 samples. This means, for example, instead of having five separate set of responses for each user, we instead have a single set of responses for each user that represents their average over the five rounds. A summary of user responses to the questions is shown in Table 4.

4.2. Path Analysis Model

To answer our three research questions, we performed Path Analysis (PA), as outlined in the Knijnenburg et al. framework for SEM [30]. Path models aim to explain how various factors are causally related to one another. In this work, we use a path model to determine which type of explanation has a stronger effect on the perceived quality of an item recommendation, and which aspects of the explanation types are driving this effect. We want to explore which explanation types are more *informative*, *entertaining* and *interesting* to participants. We create a path model that follows the Knijnenburg et al. [28] framework. In this framework, they propose grounding the model by providing the following causal paths: Objective System Aspects (OSA) $\rightarrow$ Subjective System Aspects (SSA) $\rightarrow$ User Experience (EXP). Table 3 provides the categorisation for each question asked to the participant for this framework (OSA, SSA or EXP). We created a saturated path model that captured all possible causal relations of the form OSA $\rightarrow$ SSA $\rightarrow$ EXP, and we included relations from Personal Characteristics (PC) to both SSA and EXP. We iteratively

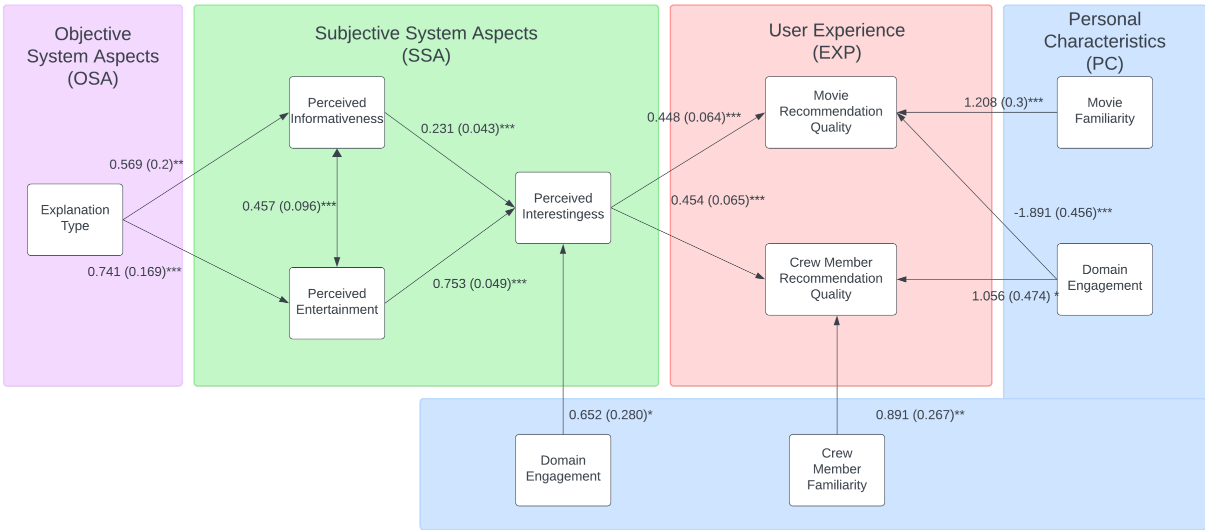

Figure 2: Path Model for explanation type with good fit (χ (17) = 24.11, p = 0.12, CFI = .981, TLI = .966, SRMR = .06, RMSEA = .06, 90% CI : [0, .116]) *: p<0.05, **: p<0.01, ***: p<0.001

Table 5
Average ratings for each explanation type on a scale from one to five. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$

Explanation Type	Mean Movie Rating	Mean Crew Member Rating	Mean Explanation Rating
Relational Explanations	3.382	2.213	2.076
Anecdotal Explanations	3.649	2.570*	2.872***

pruned non-significant paths. At each iteration, we tested each possible structural relation and removed the relation that had the highest p-value and that was not significant ($p > 0.05$), as described in [28]. The resulting path model is shown in Figure 2. This model had a good fit (χ (17) = 24.11, p = 0.12, CFI = .981, TLI = .966, SRMR = .06, RMSEA = .06, 90% CI : [0, .116]). A post-hoc power analysis on the model's non-centrality parameter (f_0) results in a power level = 0.89. In the next section, we use this model to help answer our research questions.

4.3. Research Questions

4.3.1. Effect of Explanation Type (RQ1)

To answer **RQ1**, we perform three separate two sample t-tests. We compare the mean movie ratings of the two groups of participants, and hence of the two explanation types; and we do the same for the mean crew member ratings and mean explanation ratings. In all three cases, we see (Table 5) that the means are higher for anecdotal explanations, and statistically significantly so for the crew member rating and explanation rating. Participants prefer the anecdotal explanations. But they also prefer the recommended crew member when accompanied by an anecdotal explanation. The recommender algorithm is the same in all cases. So the preference is likely to be due to the explanations. The anecdotal explanations had a more persuasive effect on participants. Participants had a higher interest in learning more about the crew IEs when presented with these anecdotal explanations. This answers **RQ1**: anecdotal explanations have a greater positive impact on perceived recommendation quality.

We note that we do not get a significant difference in mean movie ratings. This is not surprising. The explanations connect the recommended crew member to the recommended movie. Participants who answer the questions as they appear in order on the page (Figure 1) will rate the movie before seeing the explanation that connects the movie to the crew member.

4.3.2. Effect of Informative, Interesting, and Entertaining Explanations (RQ2)

In Section 4.3.1, we showed that the type of explanation has an effect on the perceived quality of the recommendations. Our path model can be used to explain what aspects of these explanations cause this effect. We explore three aspects of explanations: *informativeness*, *interestingness* and *entertainment*. From Figure 2, we can see that both movie recommendation quality and crew member recommendation quality are positively affected by the perceived interestingness of the explanation, as seen by the positive coefficients (0.448 and 0.454 respectively). Our path model also shows that an explanation's interestingness is based on how informative and entertaining it is. This is shown by the positive coefficients (0.231 and 0.753). We note that an explanation's entertainment value has a greater effect on the interestingness of the explanations. In this setting, explanations that have entertainment value have greater success than informative explanations at persuading users to consume IEs. These observations provide the answer to **RQ2**.

4.3.3. Effect of Domain Engagement and Item Familiarity (RQ3)

We use our path model to explore the effect of participants' engagement with the movie domain and their familiarity with the recommended items on their perception of the recommended items. From Figure 2, we see that participants with high domain engagement are more likely to find the explanations interesting, and the crew member recommendation to be of higher quality (coefficients 0.652 and 1.056 respectively). This makes sense, as someone who already enjoys and engages with the movie domain will be more likely to want to engage further with the domain. High engagement did lead to a lower perceived quality of movie recommendation (coefficient -1.891). With regards to item familiarity, both movie and crew member familiarity had a positive effect on their respective recommendation qualities (1.208 and 0.891). In this system, participants liked to be shown IEs that they had some degree of familiarity with. This could be due to the fact that, when the participant is familiar with an item, they feel the system is doing a good job of capturing their tastes and will respond more positively to the recommended content. It could also be due to the fact that, in the type of system we are working with (systems where users consume information about an item), participants are more open to learning about items they are already familiar with. We note that domain engagement and item familiarity had no significant moderating effect on the preferred explanation types. Participants preferred anecdotal explanations regardless of prior expertise. These observations provide the answer to **RQ3**.

5. Conclusion

We are working in the context of systems in which users consume items by reading about them. Recommender Systems play a crucial role in these systems to help users navigate and explore content that is of particular interest to them. Explanations have been shown to improve persuasiveness and user satisfaction in previous RS work. In this work, we have shown that explanations can also play this role in the kinds of systems we are working in. We use two types of explanations, *relational* and *anecdotal*, and show that the choice of explanation type can help persuade a user to learn more about an IE. We also identify three aspects of explanations, *informativeness*, *entertainment* and *interestingness*, and show how these aspects influence one another. We show that by providing explanations that are informative and entertaining, users can be persuaded to learn more about an IE. We show that a user's engagement with the domain, as well as their familiarity with the IEs, does have an impact on how open they are to exploring the domain further.

In future work, we will continue to focus on systems like MovieBuff: where users are exploring and learning about a domain. Our findings from this study will inform the design of our next generation of these systems. In particular, we will incorporate texts that seek to be informative and/or entertaining to try to persuade users to continue to explore the domain.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant No. 18/CRT/6223, which is co-funded under the European Regional Development Fund.

References

[1] F. Durao, D. Bridge, A Linked Data Browser with Recommendations, in: Procs. of the IEEE 30th International Conference on Tools with Artificial Intelligence, 2018, pp. 189–196.

[2] L. De Vocht, C. Beecks, R. Verborgh, E. Mannens, T. Seidl, R. Van De Walle, Effect of Heuristics on Serendipity in Path-based Storytelling with Linked Data, in: Procs. of Human Interface and the Management of Information: Information, Design and Interaction, 2016, pp. 238–251.

[3] A. G. Nuzzolese, V. Presutti, A. Gangemi, S. Peroni, P. Ciancarini, Aemoo: Linked Data Exploration Based on Knowledge Patterns, Semantic Web 8 (2017) 87–112.

[4] N. Tintarev, J. Masthoff, Designing and Evaluating Explanations for Recommender Systems, in: F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), Recommender Systems Handbook, Springer, 2011, pp. 479–510.

[5] M. Zanker, M. Schoberegger, An Empirical Study on he Persuasiveness of Fact-based Explanations for Recommender Systems, in: Procs. of the Joint Workshop on Interfaces and Human Decision Making in Recommender Systems, number 1253 in CEUR Workshop Proceedings, 2014, pp. 33–36.

[6] 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.

[7] C.-H. Tsai, P. Brusilovsky, Explaining Recommendations in an Interactive Hybrid Social Recommender, in: Procs. of the 24th International Conference on Intelligent User Interfaces, 2019, pp. 391–396.

[8] M. Bilgic, R. J. Mooney, Explaining Recommendations: Satisfaction vs. Promotion, in: BProceedings of Beyond Personalization 2005: A Workshop on the Next Stage of Recommender Systems Research at the 2005 International Conference on Intelligent User Interfaces, 2005.

[9] 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.

[10] B. C. G. Lee, D. Downey, K. Lo, D. S. Weld, LIMEADE: From AI Explanations to Advice Taking, ACM Transactions on Interactive Intelligent Systems 13 (2023) 1–29.

[11] V. Dominguez, P. Messina, I. Donoso-Guzmán, D. Parra, The Effect of Explanations and Algorithmic Accuracy on Visual Recommender Systems of Artistic Images, in: Procs. of the 24th International Conference on Intelligent User Interfaces, 2019, pp. 408–416.

[12] S. Gkika, G. Lekakos, The Persuasive Role of Explanations in Recommender Systems, in: Procs. of the Second International Workshop on Behavior Change Support Systems, 2014, pp. 59–68.

[13] M. Millecamp, N. N. Htun, C. Conati, K. Verbert, To Explain or Not To Explain: The Effects of Personal Characteristics When Explaining Music Recommendations, in: Procs. of the 24th International Conference on Intelligent User Interfaces, 2019, pp. 397–407.

[14] P. Kouki, J. Schaffer, J. Pujara, J. O'Donovan, L. Getoor, Personalized Explanations for Hybrid Recommender Systems, in: Procs. of the 24th International Conference on Intelligent User Interfaces, 2019, p. 379–390.

[15] N. Tintarev, J. Masthoff, Effective Explanations of Recommendations: User-centered Design, in: Procs. of the 2007 ACM Conference on Recommender Systems, 2007, pp. 153–156.

[16] J. L. Herlocker, J. A. Konstan, J. Riedl, Explaining Collaborative Filtering Recommendations, in: Procs. of the 2000 ACM Conference on Computer Supported Cooperative Work, 2000, pp. 241–250.

[17] D. Bridge, K. Dunleavy, If You Liked Herlocker et al.'s Explanations Paper, Then You Might Like

This Paper Too, in: Procs. of the Joint Workshop on Interfaces and Human Decision Making in Recommender Systems, 2014.
[18] M. Kaminskas, F. Durao, D. Bridge, Item-based Explanations for User-based Recommendations, in: Procs. of 9th International Conference on Information, Process, and Knowledge Management, 2017, pp. 65–70.
[19] P. Symeonidis, A. Nanopoulos, Y. Manolopoulos, MoviExplain: A Recommender System with Explanations, in: Procs. of the Third ACM Conference on Recommender Systems, 2009, pp. 317–320.
[20] J. Vig, S. Sen, J. Riedl, Tagsplanations: Explaining Recommendations using Tags, in: Procs. of the 14th International Conference on Intelligent User Interfaces, 2009, pp. 47–56.
[21] H. Cramer, V. Evers, S. Ramlal, M. Van Someren, L. Rutledge, N. Stash, L. Aroyo, B. Wielinga, The Effects of Transparency on Trust in and Acceptance of a Content-based Art Recommender, User Modeling and User-Adapted Interaction 18 (2008) 455–496.
[22] M. Radensky, D. Downey, K. Lo, Z. Popovic, D. S. Weld, Exploring the Role of Local and Global Explanations in Recommender Systems, in: Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems, 2022, pp. 1–7.
[23] C.-H. Tsai, P. Brusilovsky, The Effects of Controllability and Explainability in a Social Recommender System, User Modeling and User-Adapted Interaction 31 (2021) 591–627.
[24] D. C. Hernandez-Bocanegra, J. Ziegler, Conversational Review-based Explanations for Recommender Systems: Exploring Users' Query Behavior, in: Procs. of the 3rd Conference on Conversational User Interfaces, 2021, pp. 1–11.
[25] M. Zanker, The Influence of Knowledgeable Explanations on Users' Perception of a Recommender System, in: Procs. of the Sixth ACM Conference on Recommender Systems, 2012, pp. 269–272.
[26] N. Wang, H. Wang, Y. Jia, Y. Yin, Explainable Recommendation via Multi-Task Learning in Opinionated Text Data, in: Procs. of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, 2018, pp. 165–174.
[27] G. Friedrich, M. Zanker, A Taxonomy for Generating Explanations in Recommender Systems, AI Magazine 32 (2011) 90–98.
[28] B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, C. Newell, Explaining the User Experience of Recommender Systems, User Modeling and User-Adapted Interaction 22 (2012) 441–504.
[29] D. Müllensiefen, B. Gingras, J. Musil, L. Stewart, The Musicality of Non-musicians: An Index for Assessing Musical Sophistication in the General Population, PloS One 9 (2014).
[30] B. P. Knijnenburg, M. C. Willemsen, Evaluating Recommender Systems with User Experiments, in: Recommender Systems Handbook, Springer, 2015, pp. 309–352.