



Navigation by Revealing Trade-offs for Content-Based Recommendations

Linus W. Dietz^(✉), Sameera Thimbiri Palage, and Wolfgang Wörndl

Department of Informatics, Technical University of Munich, Munich, Germany
linus.dietz@tum.de

Abstract. Conversational recommender systems have been introduced to provide users the opportunity to give feedback on items in a turn-based dialog until a final recommendation is accepted. Tourism is a complex domain for recommender systems because of high cost of recommending a wrong item and often relatively few ratings to learn user preferences. In a scenario such as recommending a city to visit, conversational content-based recommendation may be advantageous, since users often struggle to specify their preferences without concrete examples. However, critiquing item features comes with challenges. Users might request item characteristics during recommendation that do not exist in reality, for example demanding very high item quality for a very low price. To tackle this problem, we present a novel conversational user interface which focuses on revealing the trade-offs of choosing one item over another. The recommendations are driven by a utility function that assesses the user's preference toward item features while learning the importance of the features to the user. This enables the system to guide the recommendation through the search space faster and accurately over prolonged interaction. We evaluated the system in an online study with 600 participants and find that our proposed paradigm leads to improved perceived accuracy and fewer conversational cycles compared to unit critiquing.

1 Introduction

Nowadays, the algorithmic side of (RSs) research has reached an impressive maturity, such that it has become virtually impossible to tell which algorithms are objectively the best [1]. However, this improvement primarily applies to traditional RSs domains, such as e-commerce, movies, and to some extent music. For recommendations in complex domains, such as tourism, the algorithmic advances of the earlier decades are of lesser value. This is because there are insufficient ratings available, the items are not so well defined in terms of their scope, and it has also been shown that users demonstrate different decision making behavior compared to purchasing physical products [2]. These challenges necessitate employing sophisticated preference elicitation strategies, and instead of collaborative filtering algorithms, recommendations are often computed with a content-based or knowledge-based paradigm. Given that traveling is a relatively rare, emotional, and high-stakes decision making scenario, RSs should provide users

with the opportunity to familiarize themselves with the items in the domain and refine their initial preferences, since users often struggle to declare their true preferences [3]. For instance, recommending which city to travel, is a very good fit for the conversational, content-based recommendation paradigm, since there are no ratings available, despite the existence of several data sets [4, Table 2].

Conversational RSs allow a directed search through the item space using some kind of dialog between the system and user [5]. Early approaches, such as FindMe [6], allow users to critique certain aspects of suggested items, whereas more sophisticated approaches allow for compound critiques [7]. Based on the observation that critiques with concrete examples can be useful [8], we are astonished that not much attention has been paid to informing users about the trade-offs involved in their critiquing choices. For example, many users would love to do a dream vacation to a buzzing city with outstanding cultural attractions, great food, a buzzy nightlife scene, favorable climate, at an affordable price tag. In reality, the combination of such features might be an empty set, thus, requiring compromising between conflicting preferences.

In this paper, we present a novel concept to navigate the item space that we call “*Navigation by Revealing Trade-offs*.” The motivation for this combination of a novel user interface and a corresponding recommendation algorithm stems from the observation that conversational RSs tend to neglect informing their users about the trade-off involved in their critiquing choices.

After surveying the related work in Sect. 2, we present the user interface in Sect. 3, and describe the recommendation algorithms in Sect. 4. We choose the destination recommendation domain, as there are suitable data sets available and it inherently requires to make trade-offs between certain aspects of the trip. The experimental setup of a large-scale user study with 600 participants is described in Sect. 5 and we present the results in Sect. 6. Finally, we conclude our findings and point out future work in Sect. 7.

2 Related Work

In this work, our application domain is recommending cities for tourist destinations. As opposed to the recommendation of hotels or point of interests [9], cities as items have no meaningful ratings, thus, the user profile and items need to be matched based on elicited preferences and features of the items. To improve the user modeling, Neidhardt et al. [10] proposed a factor analysis for tourist roles and personality traits to reveal seven tourist behavioral patterns. The authors used a set of travel-related pictures, which were assigned to each of the seven factors by experts. Since the destinations were also characterized in the feature space of the Seven Factor model [11], they could perform content-based filtering for destination recommendation. Herzog and Wörndl [12] proposed another travel RS where travel plans of multiple destinations satisfy user constraints such as budget and duration. The user modeling was done by binary indications of interest, i.e., check boxes, and the items were characterized using expert opinions and literature. Such expert-driven models are quite costly, thus, automated

approaches are preferable to scale the item characterization. Prior approaches using mainly location-based social network (LBSN) data have been successfully employed in point-of interest recommendation [13] or to characterize cities [14]. The previously proposed city characterization approach [14] is based on the distribution of its venues, where a higher amount of venue relative to the city size leads to a higher scoring. The corresponding user study also suggested that unit critiquing is a fruitful approach in the destination recommendation domain. In this work, we re-use the prototype¹ and domain model of CityRec [14] to build a conversational RS.

Critiquing is a popular approach of eliciting and refining user preferences in a conversational manner. It is usually associated with content-based filtering, although there are some research incorporating collaborative approaches [15] or even unstructured item descriptions [16]. One of the early systems, FindMe [6] introduced the concept of unit critiquing that can be seen as the start of conversational exploration of the search space in RSs research. The static unit critiquing was quite successful in several domains [6, 17], but there is opportunity to perform a smarter exploration of the item space [18]. For example, McCarthy et al. [7] proposed dynamic critiquing, to show how compound critiques can be generated dynamically, cycle-by-cycle by mining the feature patterns of the remaining products.

The evolution of dynamic compound critiques is the multi-attribute utility theory (MAUT) [19], which introduced a utility function to rank a list of multi-attribute products. Once the user selects a critique, the corresponding product is set as the current preference product in the user model and a new set of critiques is generated using a utility function. The MAUT was successfully evaluated against dynamic critiquing [7] thereby reducing the number of critiquing cycles. Chen et al. extended the MAUT-based approach and called it “preference-based organization interfaces” [20]. In their approach, the authors organized all potential critiques in a trade-off vector showing whether the features were compromised or improved in comparison to the current recommendation. That enabled them to determine useful compound critiques and successfully evaluate it using a computer configuration data set. However, we feel that such an approach is more suited for products with clear specifications, since in tourism, relative differences between the features values of items are of higher importance.

One major issue with critiquing is the divergence of the intended direction of exploration. McGinty et al. [21] studied selection strategies for recommending items in critiquing. Their Adaptive Selection approach resulted in a reduction in critiquing cycles and they could prove that their critiquing-based approaches would converge faster than preference-based approaches. Another important insight of their work was that the user should not lose the progress, i.e., the previous recommendation should be included in the upcoming cycle.

Based on these observations, we introduce a paradigm to navigate the search space that we call “*Navigation by Revealing Trade-offs*.” We propose a user interface element that visualizes the trade-offs involved in choosing one item instead of

¹ <https://github.com/myftija/cityrec-prototype>.

another in a less technical way than the preference-based organization interfaces by Chen and Pu [20]. Distinctively, our proposed interface gives the user an indication of the search space, i.e., where the current item’s feature are located within the whole feature space, which was not given in the dynamic and compound critiquing approaches [7, 22]. Furthermore, we used a utility function that determines the proposed items, aimed to resolve the “*wishful-thinking problem*” of users requesting item characteristics from the RS that do not exist in reality.

3 A User Interface Concept for Revealing Trade-Offs

3.1 Domain Model

The pure content-based paradigm requires each item to be characterized along the same features to compute recommendations. In our case, we used an available data set of already characterized 180 cities all over the world [14]. This dataset comes with a score for each city in the categories of “Food”, “Nightlife”, “Arts & entertainment”, “Outdoor and recreation”, “Cost of living”, “Shops and services”, “Average temperature”, “Average precipitation”, and “Venue count”. The domain of traveling successfully motivates our approach, since these features are natural in competition, i.e., a larger city with abundant cultural scene usually has higher cost of living, or, conversely, the nightlife options might be limited in small cities.

3.2 User Interaction

The user interaction through a web browser² goes through three major steps:

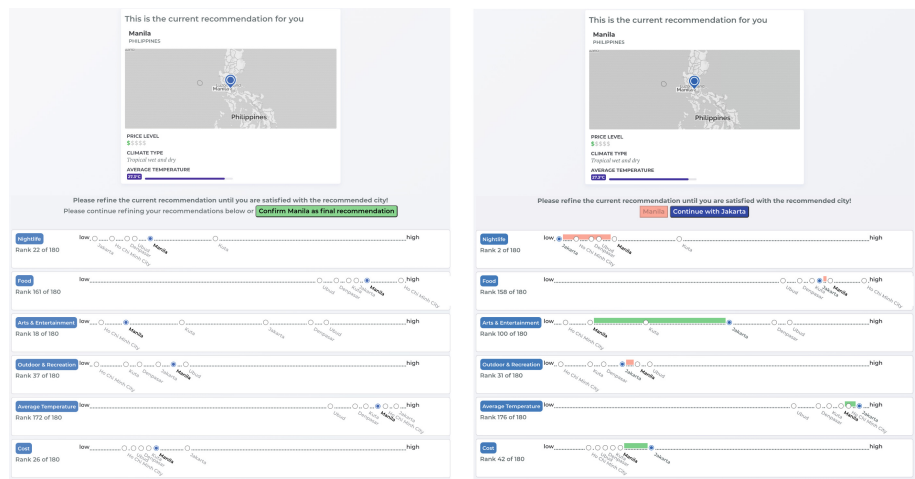
- Step (1): An initial user *Preference Elicitation Page*, where the system learns general user preferences,
- Step (2): *Conversational Refining* of the recommendations, where the user can refine preferences and learn about the trade-offs in choosing an alternative destination, and
- Step (3): *Final Recommendation Page*, where the user is shown the result.

The contribution of this paper focuses on Step (2), the *Conversational Refining*. However, this key step must be seen in the context of the whole interaction design, which we now present step-by-step.

Initial Preference Elicitation. Before the user can start refining, an initial item needs to be determined. Ideally, the system would already have an established user profile based, e.g., through previous interactions. As we have no prior information about the user, we used a previously proposed approach to present the user with an initial seed of destinations where the user can select 3–5 [14]. This seed comprises of randomly selected candidates of various clusters. By this,

² The system is available under <http://conversational-cityrec.cm.in.tum.de>.

the diversity of the sample is warranted, as the user is presented with a representative set of items to choose from. Also, this method is quite fitting for the domain and the initial set of selected cities can directly serve as input for the utility functions of Step (2). We do not aim to evaluate this method from literature [14], as we used it in the same way in all experimental conditions.



(a) Preference Refining Page. This is shown to the user at the beginning of each conversational cycle. The current city is marked bold, and five alternatives are displayed on the spectrum of each feature.

(b) Trade-off Visualization. The green and red shades indicate the trade-off involved should the user choose Jakarta instead of Manila. The user explores various alternatives before continuing.

Fig. 1. User interface of navigation by revealing trade-offs. (Color figure online)

Navigation by Revealing Trade-Offs. Figure 1 shows the interface element for our conversational “*Navigation by Revealing Trade-offs*” approach. At the top of the page, the currently recommended city is shown; below is the novel user interface. This component shows the current city along with five other cities recommended based on the utility function. For each feature the five candidate items are shown in an ordered list from low to high depending on the score. Users can select an item to see the feature value differences in all feature spaces compared to the currently recommended city. An increase in feature value is indicated using a green shade, a decrease is shown in red.

If the user is satisfied with the current recommendation, the user can choose not to continue with refining, but to confirm the current recommendation. In this case, the user is forwarded to the final recommendation page.

Final Recommendation Page. This page shows the final recommendation to the user along with a survey to measure the performance of the recommendation approaches. The final recommended city is shown with details such as the city name, country and feature values.

Baseline System. To evaluate our proposed approach, we used a modified version of the “CityRec” destination RS [14], where one could critique features of several destinations to refine them by buttons indicating “much lower”, “lower”, “just right”, “higher”, “much higher”. As the source code of this system was readily available, we used it as foundation for our experiments. We re-used the system architecture and the front-end for the initial Preference Elicitation page Step (1), and the Final Recommendation page in Step (3). However, notable differences in the user interface are that we did not use photos of cities to avoid bias due to the selection of images. Furthermore, we re-worked the unit critiquing algorithm to make it more comparable with our system. The critiques can be selected below using the same labels and logic for the adjustment as in the original approach [14], although, it is possible to adjust all features at once and the user is not limited in the number of critiquing cycles, thus, can refine the items until she is satisfied with the recommendation.

4 Algorithms

Having described the user interface elements, this section presents the machinery that computes the recommendation and, therefore, directs the path the user takes through the search space. To enable reproducibility, the system and the study data set are available under an open-source license as a Dockerized software project on Github.³

4.1 Cold-Start User Modeling

Recall that in Step (1) of the system the initial input comprises a set of 3–5 items that are characterized along the aforementioned eight features. This already allows us to compute an initial user model by simply representing the user model as an eight dimensional vector with the mean feature values of the initial cities. Nevertheless, this method is quite simple and could be interchanged with any other strategy if more information about the user’s preferences is available. Since this is not the case in our evaluation prototype, we used this simple method from literature.

4.2 Candidate Selection Strategy

The next step in the user interface requires finding candidates of which the user can choose one to progress the search for suitable recommendations. In typical content-based recommendation style, one could naively use any similarity metric, such as the Euclidean Distance on normalized feature values to compute some cities similar to the current user model. The top items can then be shown as alternatives to the user. One issue with this strategy is, that it does not consider the user preference variations during the refinement. Furthermore,

³ <https://github.com/LinusDietz/conversational-cityrec>.

the convergence of the algorithm will be poor, since it presents the user with similar items to the current recommendation, thus, the user will not have the option to select a city with a significantly different feature value. Instead, we propose the “Variance Bi-distribution” utility function (Eq. 3) whose value is defined by two normal distributions per feature, each representing an increase or decrease of feature value. The two normal distributions are given as $\sim N(\mu_1, \sigma^2)$ and $\sim N(\mu_2, \sigma^2)$, where μ_1 and μ_2 define the position of the bell curves on the normalized value range of the feature, and σ defines the shape of the curve.

The distance between the currently selected reference item, ref_k and the respective bell curves are computed by adding or subtracting an offset computed in Eq. 1. This offset is the standard deviation of each feature value f of all previous items in the conversational history H by the number of previous conversational iterations n . The numerator of the offset needs to be moderated by a constant C_m , which we empirically determine for the dataset in Sect. 5.1. To summarize, the mean of the normal distribution is farther from the current user model if the variance of a feature is higher.

$$\mu_1 = ref_k - \frac{\sqrt{Var(f \in H)} \cdot C_m}{n} \quad \mu_2 = ref_k + \frac{\sqrt{Var(f \in H)} \cdot C_m}{n} \quad (1)$$

The second parameter of the normal distributions, σ , is computed in a similar way (cf. Eq. 2). This has the effect that with a higher variance, we obtain a flatter distribution and, thus, a lower impact of this feature on the utility score.

$$\sigma = \frac{\sqrt{Var(f \in H)} \cdot C_s}{n} \quad (2)$$

The intuition behind this is that if the user has a strong preference regarding a feature having a certain value and consistently picks cities with a high temperature, the system is quite certain of this user’s preference toward temperature and, thus, should put high weight to this feature. Conversely, if a user has selected cities with another feature having both low and high values resulting in a high variance, it can be seen as a signal that the user has no specific preferences toward the feature as it is not of importance to the user. Thus, the impact of such a high-variance feature should be smaller than a low-variance feature. Over time, we increase this effect by dividing through the number of previous iterations n . This further helps the algorithm converge.

The maximum score of the two distribution functions for a given item feature is taken as the utility score of the respective feature. We then compute the overall utility of each item as the sum of all feature scores of the utility function.

$$utility = \sum_{f \in F} s(f) \quad (3)$$

Convergence Behavior. The effect of this utility function is that it balances exploration in the beginning and fine-tuning in later stages of the search. If

a feature variance is high and the number of iterations small model adjusts μ_1 and μ_2 further away from the reference point, with a higher σ resulting in a flatter distribution of the feature’s utility function. In this case, items far away in the feature space also would get higher utility scores, ensuring users are presented with cities more spread across the feature space. With a larger number of iterations, the user preferences for particular features are converging, i.e., the user will be presented with an increasingly narrower band of feature values to refine the preferences. As a result, μ_1 and μ_2 are closer to the feature value of the current recommended item, with a smaller σ , such that items with similar feature values have a substantially higher utility score than the cities with dissimilar feature values. However, if the variance of a feature is still high, the curve will stay quite flat giving this feature less weight, thus, recognizing that the user is rather indifferent toward this feature. This convergence behavior can be observed in Fig. 1. After some iterations, the algorithm determined that the user has a clear preference for high scores in the food and temperature aspects, and low scores in nightlife, outdoor & recreation, and cost. Thus, the refining candidates are quite close by each other, whereas they are spread along the spectrum in the arts & entertainment spectrum.

Elimination of Candidates. To further improve the convergence, we propose a variant that eliminates items whose feature values have been refined in a contrary way. The reasoning behind this elimination of candidates is that if a user refines a feature of an item, it becomes an explicit information that the value of the feature is unsatisfactory and should take only values toward the direction of the refinement. Thus, we can compute candidates just as before, however, items that have a lower (or higher) value than the original item ref_k are removed from the search space. For example, if the user refines the value of *Arts & Entertainment* of Manila in Fig. 1b in favor of Jakarta, the system will assume that all cities that have a lower value in *Arts & Entertainment* than Manila should be excluded from future suggestions.

5 User-Centric Evaluation

For the evaluation of the system, we chose a between-subject design to perform a large-scale online user study. First, we need to determine the constants of the Variance Bi-distribution Model for the current data set.

5.1 Instantiation for the Domain

During the development of the system, we noticed that using only the standard deviation divided by the number of iterations in Eq. 1 and 2, μ_1 and μ_2 would be too extreme, which will result in items recommended that are too far away from the current city. To moderate this effect, the constants C_m and C_s of Eqs. 1 and 2 were introduced for the Variance Bi-distribution Model. This step ensures an efficient navigation should be seen as an adjustment of the algorithmic properties to the data set at hand, as different domains can have different characteristics, i.e., a different number of items.

Determining Constants. The values of C_m and C_s can be determined in an offline setting using a simulation. This is because by systematically altering the values of C_m and C_s , we can see how quickly the algorithms converge from an initial setting after Step (1) to a desired item while making *consistent decisions*. In the context of the simulation, we define consistent decisions by choosing the item that is nearest to the target item using the distance metric of the RS. Thus, the simulator chooses candidates toward the target recommendation, just as a real user would, until that recommendation is part of the set of candidate items. For the cities, we used user interaction data to perform a realistic simulation [14]. The data set of 63 user sessions contained the initial city selections by the user and the final recommendation the user had selected. Having historic data for the simulation, we can now train the parameters using relevant scenarios, as opposed to the randomized or exhaustive simulation strategies.

Result. Regarding parameters of the simulation, we varied C_m from 2 to 6, and C_s from 4 to 20, both in 0.5 intervals. For each these parameters' configuration, we recorded the session length of the 63 user sessions of the data set. The result of the simulation reveals a global optimum at $C_m = 3$ and $C_s = 8$.

5.2 Online User Study

We conducted the user study using the online experimentation platform Prolific.⁴ We used a between-subject design and invited participants of the platform who had indicated “Traveling” as one of their hobbies. Only one independent variable was randomly assigned to the users, i.e., the critiquing system in Step (2). The three options⁵ were the baseline unit critiquing system and the trade-offs UI using the Variance Bi-distribution Model without and with the elimination variant. As dependent variables, we used metrics about the user interaction and a subset of the ResQue Questionnaire (cf. Table 1), which is a validated, user-centric evaluation framework for RSs [23], where users indicate their agreement with each statement on a Five-point Likert Scale.

6 Results

The user study was conducted in December 2020 with 600 participants. Out of the 600 participants, we excluded 181 responses, which failed an attention check, showed very low interaction with the system, i.e., an interaction of less than 35s, and did not use a desktop browser as instructed. This left us with 419 valid submissions (59.9% female, 39.1% male, 1% other) from 42 different countries. The users predominantly came from Europe, due to the time zone when the survey was initiated. The age distribution was 20.8% of below 21 year olds, 55.6% were 21–30, 13.1% were 31–40, 6.2% were 41–50, 3.1% were 51–60,

⁴ <https://prolific.co/>.

⁵ The variants can be tested under <http://conversational-cityrec.cm.in.tum.de>.

Table 1. Hypothesis testing of the dependent variables between the baseline unit critiquing and the two variants of the Trade-off refinement. The mean values of the survey items coded as integers from 1 to 5 are for informative purposes only.

Variable	Baseline	Trade-offs			Trade-offs w. Elim.		
	Mean	Mean	p	w	Mean	p	w
(Q1) Interest match	3.81	4.12	0.002	7378	4.07	0.005	8727.5
(Q2) Better than friend	3.26	3.25	0.939	9053.5	3.26	0.749	10212.5
(Q3) Cities are familiar	4.09	4.14	0.605	8794.5	4.22	0.187	9572
(Q4) Rec. cities are attractive	4.06	4.18	0.314	8524	4.05	0.538	10816.5
(Q5) Discover new Cities	3.66	3.76	0.42	8608	3.71	0.711	10179
(Q6) Adequate layout	3.78	3.45	0.003	10917.5	3.56	0.044	11765
(Q7) Easy to modify preferences	4.14	3.59	<0.001	11846.5	3.66	<0.001	13235
(Q8) Became familiar quickly	4.19	3.67	<0.001	11681	3.60	<0.001	14125
(Q9) Influenced selection	3.44	3.64	0.043	9104.5	3.63	0.044	9104.5
(Q10) Overall satisfaction	3.82	3.84	0.743	8910.5	3.77	0.534	10824.5
Number of conversational cycles	4.44	2.38	<0.001	–	2.46	<0.001	–

and 1.2% were 61 years or older. With respect to the independent variables, 140 were assigned to the baseline unit critiquing, 130 to the Trade-off Refinement, and 149 to the Elimination Variant.

Quantitative Analysis. Regarding the number of conversational cycles, we observed that all sessions using the Trade-off interface were finished by the users within 6 cycles, with a mean value of 2.38/2.46, whereas the baseline unit critiquing interface needed more cycles with a mean value of 4.44 cycles. Thus, the Trade-off UI reduced the iterations by of 46.4% (44.6% in the elimination variant), which is a significant reduction when testing the hypothesis using a t-test (cf. last row of Table 1). Note that the user interface was set up in a way, so that at least one interaction cycle had to be performed, before the users could accept the current recommendation as final result.

For the survey items, we computed cross-wise Wilcoxon rank sum tests for independent populations using the three independent variables. The null hypotheses were that there is no difference in the median of the responses. Since we could not find significant differences between the Trade-off refining and Trade-off refining with the Elimination variant, we only tabulated the outcomes in Table 1 with respect to the baseline unit critiquing. Besides the analysis of the number of conversational cycles, we could refute the null hypothesis in favor of the Trade-off Variants in (Q1) and (Q9), while the baseline received better responses in (Q6), (Q7), and (Q8). This mixed result can be summarized in a way, that the Trade-off interface had superior perceived recommendation accuracy at the expense of the users' perceived ease of use.

Discussion. The superior perceived accuracy measured by (Q1) at about 45% fewer conversational cycles, underlines the merit of our proposed user interface. However, the subjects rated the usability-related metrics of the unit critiquing system higher (Q6–Q8). We suspect that this is due to that unit critiquing has already been employed in various RSs, so it is quite possible that many users were already familiar with this concept. Dealing with a new refinement interface involving reasoning about trade-offs certainly involves more cognitive effort and, thus, might need more familiarization (Q8) than only one session. The study was designed in a way that users could only submit the survey once and we did not familiarize the users with the system before their session to avoid learning effects. The significant difference in (Q9) “This recommender system influenced my selection of cities.” in favor of the Trade-off interface is likely an artifact of the comparative lengthy search in the unit critiquing, since both values are in the center of the Likert Scale. Interestingly, there were no significant differences in any dependent variables between the Trade-off refinement and its Elimination variant. We attribute this to the low number of conversational cycles that were needed to come up with a satisfactory result. In the given data set of 180 cities, the elimination of candidates was probably not necessary, as the utility function was able to recommend attractive items after two or three cycles. Nevertheless, we are confident that the concept of elimination of parts of the search space based on the users’ choices could be useful and we plan to analyze the merit of the Elimination variant with larger item sets of over 1000 items.

7 Conclusions

The success of modern recommender systems depends on the seamless integration of algorithms and user interface elements. Given that existing critiquing systems have often neglected to explicitly inform users about the trade-offs of the critiquing actions, we developed the Navigation by Revealing Trade-offs system, which integrates a user interface concept with a utility function to compute refinement candidates. The evaluation shows that perceived accuracy is better than the unit critiquing baseline at similar reductions in the number of conversational cycles as other advanced critiquing approaches have demonstrated [19, 21].

Based on this promising result, further analyses of this refinement paradigm should follow with larger item sets to analyze the merits of the Elimination variant. Since our study followed a between-subject design, we also can not answer whether the higher ratings for the interface adequacy are due to that unit critiquing being conceptually easier to understand or users are more familiar with such a long-established paradigm. Therefore, the usability and learnability should be investigated in a usability analysis in a controlled laboratory setting.

References

1. Dacrema MF, Cremonesi P, Jannach D (2019) Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In: 13th ACM conference on recommender systems, September 2019. ACM, New York, pp 101–109
2. Werthner H, Ricci F (2004) E-commerce and tourism. *CACM* 47(12):101–105
3. Sertkan M, Neidhardt J, Werthner H (2020) Eliciting touristic profiles: a user study on picture collections. In: 28th ACM conference on user modeling, adaptation and personalization. ACM, New York, pp 230–38
4. Jannach D, Manzoor A, Cai W, Chen L (2021) A survey on conversational recommender systems. *ACM Comput Surv* 54(5):1–36
5. Salamó M, Reilly J, McGinty L, Smyth B (2005) Knowledge discovery from user preferences in conversational recommendation. In: Jorge AM, Torgo L, Brazdil P, Camacho R, Gama J (eds) *Knowledge discovery in databases*. Springer, Heidelberg, pp 228–239. https://doi.org/10.1007/11564126_25
6. Burke RD, Hammond KJ, Young BC (1997) The FindMe approach to assisted browsing. *IEEE Expert* 12(4):32–40
7. McCarthy K, Reilly J, McGinty L, Smyth B (2004) On the dynamic generation of compound critiques in conversational recommender systems. In: De Bra PME, Nejdl W (eds) *International conference on adaptive hypermedia and adaptive web-based systems*. Springer, Heidelberg, pp 176–184. https://doi.org/10.1007/978-3-540-27780-4_21
8. Viappiani P, Faltings B, Pearl P (2006) Preference-based search using example-critiquing with suggestions. *JAIR* 27(1):465–503
9. David M, Ricci F (2018) Clustering users’ POIs visit trajectories for next-POI recommendation. In: Pesonen J, Neidhardt J (eds) *Information and communication technologies in tourism*, December 2018. Springer, Cham, pp 3–14
10. Neidhardt J, Seyfang L, Schuster R, Werthner H (2015) A picture-based approach to recommender systems. *JITT* 15(1):49–69
11. Sertkan M, Neidhardt J, Werthner H (2019) What is the “personality” of a tourism destination? *JITT* 21(1):105–133
12. Herzog D, Wörndl W (2014) A travel recommender system for combining multiple travel regions to a composite trip. In: *CBRecSys*, pp 42–48
13. Zhiwen Yu, Huang X, Yang Z, Guo B (2015) Personalized travel package with multi-point-of-interest recommendation based on crowdsourced user footprints. *IEEE Trans Hum-Mach Syst* 46(1):151–158
14. Dietz LW, Myftija S, Wörndl W (2019) Designing a conversational travel recommender system based on data-driven destination characterization. In: *ACM Rec-Tour*, September 2019, pp 17–21
15. Xie H et al (2018) Incorporating user experience into critiquing-based recommender systems: a collaborative approach based on compound critiquing. *Mach Learn Cybern* 9(5):837–852. <https://doi.org/10.1007/s13042-016-0611-2>
16. Rana A, Bridge D (2020) Navigation-by-preference: a new conversational recommender with preference-based feedback. In: 25th International conference on intelligent user interfaces, March 2020. ACM, New York, pp 155–165
17. Burke RD (2002) Interactive critiquing for catalog navigation in e-commerce. *Artif Intell Rev* 18(3):245–267
18. McGinty L, Reilly J (2011) On the evolution of critiquing recommenders. In: Ricci F, Rokach L, Shapira B, Kantor P (eds) *RecSys handbook*, October 2011. Springer, Boston, pp 419–453. https://doi.org/10.1007/978-0-387-85820-3_13

19. Zhang J, Pu P (2006) A comparative study of compound critique generation in conversational recommender systems. In: Wade VP, Ashman H, Smyth B (eds) Adaptive hypermedia and adaptive web-based systems. Springer, Heidelberg, pp 234–243. https://doi.org/10.1007/11768012_25
20. Chen L, Pu P (2007) Preference-based organization interfaces: aiding user critiques in recommender systems. In: Conati C, McCoy K, Paliouras G (eds) User modeling. Springer, Heidelberg, pp 77–86. https://doi.org/10.1007/978-3-540-73078-1_11
21. McGinty L, Smyth B (2006) Adaptive selection: an analysis of critiquing and preference-based feedback in conversational recommender systems. *Electron Commer* 11(2):35–57
22. Reilly J, Zhang J, McGinty L, Pu P, Smyth B (2007) Evaluating compound critiquing recommenders: a real-user study. In: 8th ACM conference on electronic commerce. ACM, pp 114–123
23. Pu P, Chen L, Hu R (2011) A user-centric evaluation framework for recommender systems. In: 5th ACM conference on recommender systems. ACM, New York, pp 157–164

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

