Designing Interactive Visualizations of Personalized Review Data for a Hotel Recommender System

Catalin-Mihai Barbu
University of Duisburg-Essen
Duisburg, Germany
catalin.barbu@uni-due.de

Jürgen Ziegler
University of Duisburg-Essen
Duisburg, Germany
juergen.ziegler@uni-due.de

ABSTRACT

Online reviews extracted from social media are being used increasingly in recommender systems, typically to enhance prediction accuracy. A somewhat less studied avenue of research aims to investigate the underlying relationships that arise between users, items, and the topics mentioned in reviews. Identifying these—often implicit—relationships could be beneficial for at least a couple of reasons. First, they would allow recommender systems to personalize reviews based on a combination of both topic and user similarity. Second, they can facilitate the development of novel interactive visualizations that complement and help explain recommendations even further. In this paper, we report on our ongoing work to personalize user reviews and visualize them in an interactive manner, using hotel recommending as an example domain. We also discuss several possible interactive mechanisms and consider their potential benefits towards increasing users' satisfaction.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; *Personalization*; • **Human-centered computing** → *User centered design*;

KEYWORDS

Recommender systems, Personalized reviews, Interactive visualization, Tourism, Multimode networks, Trustworthiness

1 INTRODUCTION AND MOTIVATION

As the research focus in recommender systems (RS) shifts gradually from prediction accuracy towards more user-centric methods, topics such as personalizing the user experience and increasing users' trust become more salient [11]. Transparency [18] and control [6] are frequently mentioned in the literature as important factors for achieving these goals. In this context, various approaches have been developed to support users in their exploration of recommendations. Collectively, these approaches are usually referred to as *interactive recommending* [7].

When many attributes need to be considered before making a choice, as is often the case in hotel RS, comparing ranked lists of recommendations often becomes cumbersome [3]. At the same time, alternative visualization techniques need to strike a fine balance with respect to the amount of information that can be presented while maintaining ease of understanding. Because of this inherent difficulty, ranked lists are frequently, despite their shortcomings, the preferred way to display recommendations. A promising middle-ground approach is to visualize specific aspects of a recommendation (e.g., user-generated content) while still retaining the traditional presentation style for the item lists. Prior research has established that online reviews can be a rich source of contextual information [2, 4, 26]. When presented alongside factual product attributes and standardized ratings, reviews can provide additional background evidence to support users in their decision-making process. Consequently, reviews are being used—with increasing effectiveness—as a further means of explaining recommendations [4, 20]. At the same time, large amounts of user-generated content also create an opportunity for personalization.

In this paper, we describe our ongoing approaches to personalize user reviews for a hotel RS and to visualize them in an interactive manner. The contribution of our work is threefold, namely to: 1) propose a model for identifying a suitable set of reviews to show a specific user, taking advantage of implicit relationships mined from those reviews; 2) develop methods to visualize said reviews to support users' decision-making; and 3) explore interactive mechanisms that allow users to maintain control over the visualization. Our approach builds upon the *co-staying* concept introduced in [1], wherein implicit multimode (user-topic-item) relationships extracted from user-generated content may be useful for increasing the trustworthiness of hotel recommendations.

In the following section, we report on the state of the art in review personalization and in information visualization techniques for RS. Afterwards, we present our conceptual model for personalizing reviews, using hotel recommendations as an example domain. Subsequently, we propose an approach for visualizing the data based on Sankey diagrams [24]. We also describe several mechanisms for interacting with the visualizations. Finally, we conclude by reflecting on our approach and enumerating promising directions for future research.

2 RELATED WORK

Although the importance of online reviews for explaining recommendations has been recognized in prior work (see, e.g., [20] for an overview), the topic of personalizing the presentation of reviews in RS has received relatively little attention from researchers. Moghaddam et al. [14] provides empirical evidence to support the fact that the perceived quality and helpfulness of online reviews differs across users. Their evaluation, which was performed on a real-life dataset of reviews, compared two latent factor models for predicting the personalized review quality. Similarly, Tu et al. [21] aim to personalize the set of reviews shown to users by decreasing redundancy and maximizing the coverage of topics of interest. Once a suitable set of reviews has been identified, the next challenge is how to present them.

Information visualization for RS is an active and promising field of research [10]. Several approaches have been proposed for visualizing recommendations in an interactive manner. We believe some of these approaches could also be adapted for visualizing specific aspects of a recommendation. In *SetFusion* [16], a hybrid RS for conference talks, the authors enhance the typical ranked list paradigm with interactive Venn diagrams. The charts afford users a new perspective on examining and filtering recommendations. The implementation is a successor of *TalkExplorer* [23], in which the relevant information was represented using cluster maps. Yazdi et al. [25] propose a bubble graph representation for suggesting collaboration opportunities. They show that the visualization helps users form a mental model of the recommendation space and the connections between scholars, institutions, and research topics. A similar visualization metaphor is also used in [15] to recommend contacts in social networks. Richthammer and Pernul [17] employ treemapping to facilitate users' exploration of movie recommendations. They show that the structured presentation makes it easier for users to obtain an overview of the search space and possible alternatives. In contrast, Kunkel et al. [12] render the movie domain space on a 3D map that can be reshaped by users to "uncover" similar recommendations. Finally, Tietz et al. [19] proposes a method for visualizing multimedia content based on linked data. Displaying the semantic relationships graphically supports exploration and the discovery of new content.

Despite the multitude of techniques, most of them are inherently limited in the number of elements that can be realistically depicted on a screen. Thus, identifying and grouping items into clusters becomes a key requirement for reducing clutter and helping users cope with the amount of information. Several approaches have been proposed in the field of social network analysis that can be applied to multimode networks (see, for instance, [8], [9], and [13]).

3 PERSONALIZE A SET OF HOTEL REVIEWS

Whether a hotel review is considered helpful by a user may depend on several aspects, among them individual preferences (e.g., "I prefer to sleep on a soft mattress; what have previous guests written concerning bed quality?"), the specifics or requirements of the travel scenario (e.g., "I am traveling for work, so I am mostly interested in the opinions of other business travelers."), and various sociodemographic factors (e.g., "What do people who, like me, usually book 3-star hotels think about these accommodations?"). The goal of personalization is to show users the most relevant reviews, based on their recorded preferences [21]. Our hypothesis is that both the content of the review and metadata about the person who wrote it can be leveraged to calculate a relevance score. This would allow a RS to prioritize hotel reviews that: 1) mention the topics in which the user is interested; and, at the same time, 2) are written by people who have the most in common with the user.

Various techniques have been proposed for extracting features and user attitudes from online reviews [2, 4, 26]. Most commonly, the output is a list of concepts, or topics, that appear often in reviews (for example, "soft bed" or "quiet room"). User attitudes about a certain topic can be either positive, negative, or neutral [26]. In [1], we described how the connections between users, hotels, and topics form an *implicit* social network—meaning that users do not communicate directly with each other. Instead, relationships are formed based on the hotels that they have visited in the past and the topics that they mentioned in their reviews.

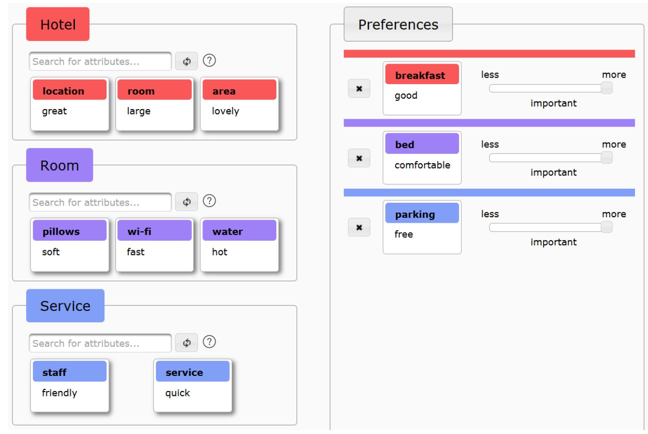

Figure 1: Eliciting user preferences. Users can drag and drop relevant topics from the categories on the left-hand side to the "Preferences" area on the right-hand side. Sliders can be used to adjust the importance of each attribute.

For the sake of simplicity, and as an initial step towards testing our hypothesis, we decided to elicit user preferences as part of the recommendation process. Concretely, in our application—which is based on the one described in [5]—users shall be asked to select (and assign weights to) hotel characteristics that are most important to them (Figure 1). This interaction bears similarities to how a person typically interacts with online booking portals: After choosing a destination and travel date, users are normally presented with a list of filters that they may use to refine the list of recommendations even further. Clicking on a filter labeled "beach", for instance, will prioritize hotels located near the seafront. Such an action can be regarded as preference elicitation. In our prototype, we feed this information into the RS not only to find recommendations, but also to personalize the reviews.

Once they have been elicited, user preferences can be matched against pre-extracted topics (see [5]) to select the most suitable reviews. For each review belonging to one of the recommended hotels, a partial relevance score, R_c, can be computed based on the number of topics that match the user preferences. A second, and arguably more interesting step, is to additionally consider user similarity when calculating a review's final relevance score. We identified four user factors that we consider relevant for this task. A reviewer's *rating behavior* denotes the extent to which her hotel scores match those of other users who share similar preferences. This is, in essence, the basis for collaborative filtering [11]: For a given set of hotels, we expect like-minded guests to give more homogeneous ratings. The *travel profile* represents a combination of aspects that characterize the reviewer's typical hotel booking. These may include the purpose of travel (i.e. business or leisure), room type, number of nights, time of year etc. Another factor is the degree to which a reviewer's own *set of preferences* is well-defined. For example, reviews contributed by someone who often gives feedback on the quality of the bed are probably more relevant to a user who cares about this aspect of a hotel room. Finally, we check whether the reviewer has stayed in *similar hotels*. For this, we consider both objective information, such as a hotel's star rating, and prevalent topics extracted from user-generated content. Prior work suggests

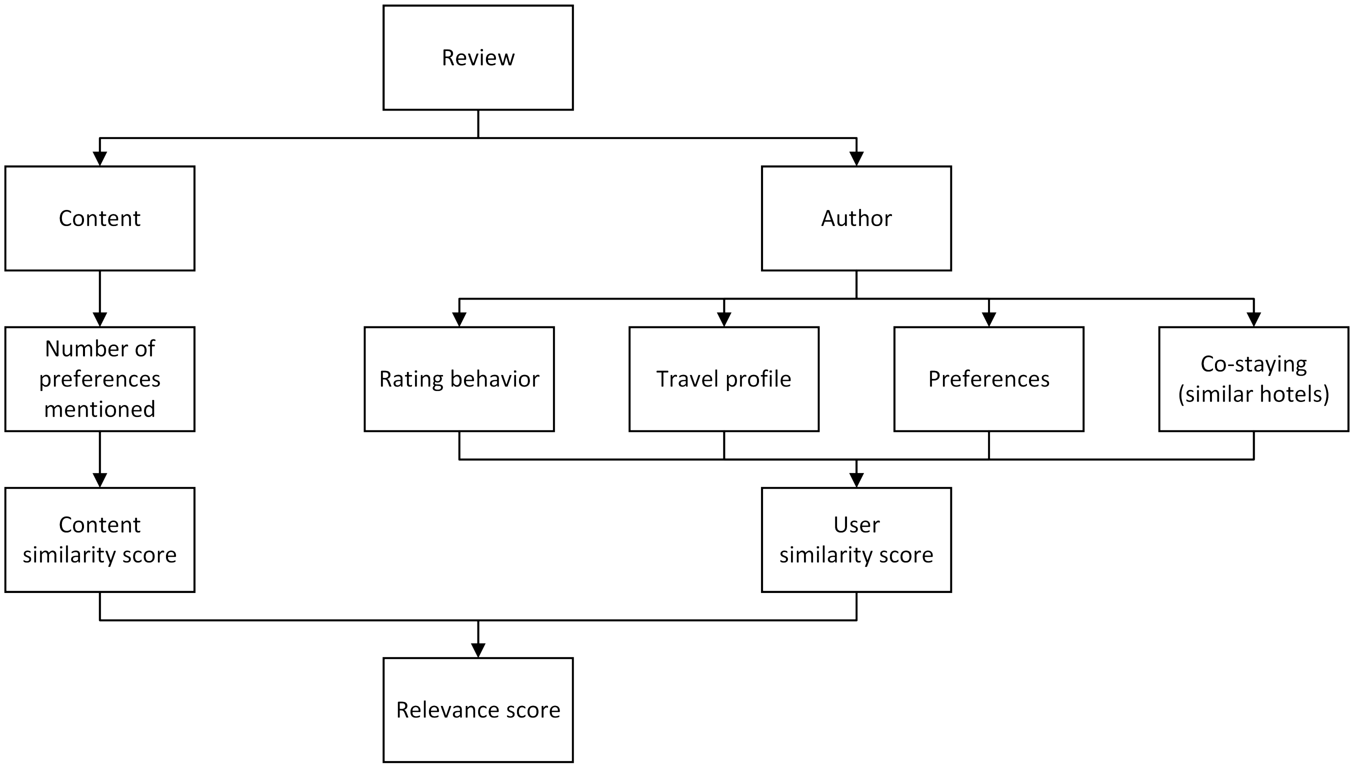

Figure 2: Proposed model for calculating the relevance score of a review, taking into account both its content and its author.

that people who book similar hotels may also have comparable expectations [2, 21]. By combining these factors, the second partial relevance score, R_{u}, can be calculated. The review's final relevance score can be written as $R_{\mathrm{r}} = R_{\mathrm{c}} \cdot wc + R_{\mathrm{u}} \cdot (1 - wc)$, where the weighting factor wc will be found empirically. An overview of the proposed model is shown in Figure 2. With the exception of travel profile, all factors can be extracted from information contained in the co-staying network. The remaining factor can be obtained from the reviews' metadata (i.e. the review date and automatically-generated tags about the hotel booking, such as duration of stay, type of room, and number of guests). The dataset used for generating the co-staying network is the one described in [1].

As a further refinement, we will also explore the possibility of using reviews written for hotels that are part of the same chain as the recommended hotels. Our premise is that hotel chains typically strive to achieve a consistent user experience across their sites [1]. This means that, per our *co-staying* concept, two reviewers can be considered similar even if they previously booked rooms in different locations of the same hotel franchise. We aim to evaluate our approach by comparing it against latent factor models, such as the one suggested in [14]. We believe the additional relationships captured by the multimode network will yield improved results when compared to other review personalization approaches.

4 VISUALIZE AGGREGATED REVIEW DATA

Based on our review of the literature (see section 2), we believe there is significant potential in combining traditional RS with a means to explore information related to a specific hotel recommendation in a more visual manner. Concretely, we started developing graphical representations of relevant hotel topics (and their authors) based on: 1) how often they appear in the user-generated content; and 2) their valence (i.e. positive or negative mentions). To avoid information overload, we purposefully restrict the visualization to only a personalized set of hotel reviews, as identified in the previous section. Our aim is to find out whether such a visualization has a significant effect in terms of helping users understand better why a hotel was recommended. Thus, we consider the visualization as an additional form of explanation. Constraining the visual representation to relatively small amounts of data (i.e. from a personalized subset of reviews) also alleviates the main shortcoming identified in the related work section. At the same time, we believe our approach remains in line with the typical use cases of hotel RS. Specifically, most people have a limited number of preferences (i.e. topics) in which they are interested in for a given trip.

We experimented with two graphing methods, namely: 1) Treemap, an area-based visualization [17]; and 2) Sankey, a type of flow diagram [24]. Both techniques have specific advantages and shortcomings. In general, Treemaps provide a good overview, but users might find it more difficult to focus on specific details. In contrast, Sankey diagrams tend to have a higher legibility. This is due to their flow structure, which generally follows a left-to-right (or, less frequently, top-to-bottom) orientation that might be easier for users to grasp. Because of this aspect, we will focus on Sankey visualizations in the remainder of this paper.

The layout of a Sankey diagram is flexible enough to accommodate multiple levels of nodes. As a result, it is well-suited for

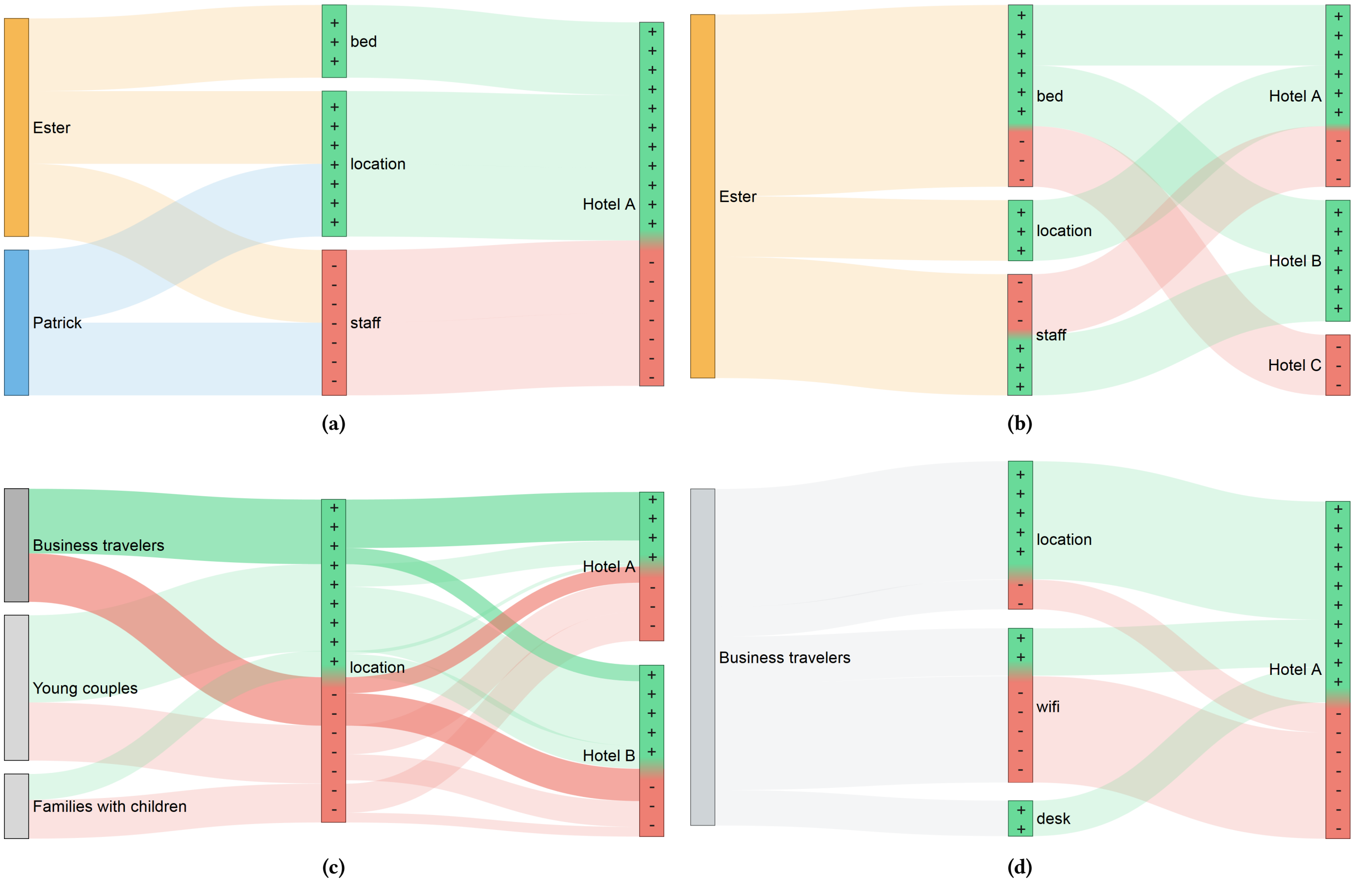

Figure 3: Example visualizations using Sankey diagrams. Different colors (green and red) and symbols ("+" and "-") are used to denote positive and negative mentions, respectively. Top left: Topics mentioned by two users in their reviews about Hotel A. Top right: Topics mentioned by a user in her hotel reviews. Bottom left: Opinions regarding the location of two hotels have been aggregated based on users' travel category. Links originating from the group "business travelers" are highlighted. Bottom right: Subset of topics mentioned by a group of users who reviewed Hotel A.

visualizing multidimensional data, such as the user-topic-hotel relationships that form the backbone of our co-staying network. Four typical visualizations are shown in Figure 3. Each follows a similar pattern, with the user (or user group) nodes placed on the left, topic nodes in the middle, and hotel nodes on the right. Edges between nodes correspond to topic mentions; the width of an edge is proportional to the number of times its corresponding topic appears in a user's reviews. User sentiment is represented using colors (i.e. red and green for negative and positive mentions, respectively) and symbols (i.e. "-" for negative and "+" for positive mentions). Furthermore, the coloring of topic and hotel nodes indicates the proportion of positive vs. negative references. These graphical elements are meant to help users perceive quickly the prevailing user sentiment on a given issue. Specific paths in the Sankey diagram can be highlighted to increase their salience, as shown in Figure 3c. As depicted in Figure 3c and Figure 3d, the visualization can also be used to compare two or more hotels.

Since many prospective users might not be familiar with Sankey diagrams, we formulate several interactive mechanisms to support them. First, and most importantly, users should be able to control the amount of information that is represented in the chart. One way to achieve this is by clustering nodes to reduce clutter and increase legibility. This is especially relevant in the case of user nodes, which will almost always be the most numerous of the three vertex types. A relatively straightforward possibility is to group users based on whether they are traveling for business or leisure (Figure 3c). A more interesting approach that we are investigating is how to cluster users based on their similarity scores, which are computed using the algorithm described in the previous section. Furthermore, topics can also be clustered, for example based on whether they refer to the hotel in general (e.g., "location"), a room feature (e.g., "shower"), or the quality of the service (e.g., "staff").

Users will also have the option to "zoom" in or out in order to fine tune the level of detail. Another way to control the visualization is by providing adequate filtering mechanisms. For example, the user may select only a subset of topics to visualize, or she might decide to view only topics with negative opinions. Even so, showing all three layers of the underlying multimode network at once might still prove too difficult for some users to comprehend. Therefore, one possible solution is to limit the visualization to only two types of

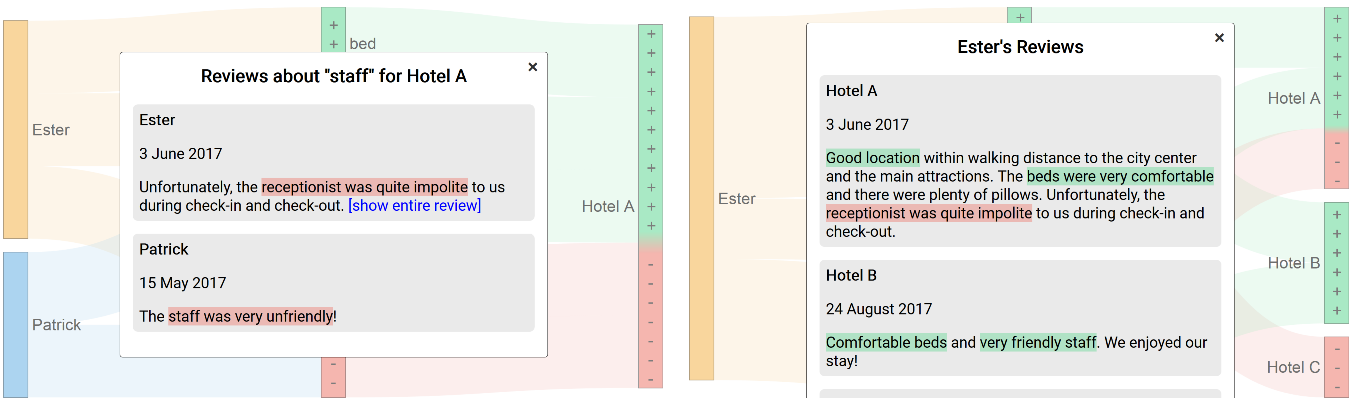

(a) (b)

Figure 4: Reviews are displayed (on demand) as a separate layer on top of the Sankey visualization. Topics are highlighted according to their valence. Left: Users' opinions about a particular topic related to Hotel A. For the top review in the list, only a relevant snippet is shown. Right: Partial view of the reviews written by a user.

vertices. In this case, suitable interface elements could be provided to facilitate interaction with the third dimension, e.g., by using filters.

Clicking on the nodes also affords interesting interaction opportunities. One example is to allow users to "refocus" the visualization around a specific node. In Figure 3a, clicking on one of the two users changes the diagram to show only the topics mentioned by that user (Figure 3b). Similarly, selecting a topic would display only the users who referred to that topic in their reviews. Finally, clicking on the hotel would have the effect of reverting to the default visualization. An interesting open question, which we plan to verify empirically, is whether to allow users to reorganize the diagram by dragging and dropping nodes. Such functionality may facilitate "ad-hoc" clustering. Moreover, the resulting arrangement could also be saved as a template, so that future visualizations are rendered, by default, in a similar fashion.

Initially, our Sankey diagram implementation does not display the actual content of the reviews. However, users can easily access this information on demand (cf. Figure 4). One relatively simple method to achieve this functionality is to render the appropriate reviews in an overlay window. The content and presentation style are determined by the node or edge with which the user interacted. In Figure 4a, interacting with the node "staff"—e.g., by double-clicking—displays users' feedback on that topic. (Note that the underlying Sankey diagram is identical to the one in Figure 3a.) Furthermore, the top review in the aforementioned example has been condensed to a relevant snippet; however, the user may toggle an embedded link to view the entire text. By the same token, interacting with either a hotel or with a user node depicts all hotel reviews, or the opinions contributed by a specific user, respectively. An example of the latter is shown in Figure 4b (see also Figure 3b for the initial visualization). Moreover, this type of interaction is implemented for edges as well. Alternatively, a user may only be interested in finding out quickly how many times a topic has been mentioned, without perusing the reviews. In this case, simply hovering over an edge will display this information in a summarized form, e.g., "'breakfast' → 5 mentions (mostly positive)".

5 DISCUSSION AND FUTURE WORK

As the amount of user-generated content continues to grow, it is becoming increasingly important to develop methods for filtering and personalizing the content used for explaining recommendations. We propose a model for identifying personalized sets of reviews in a hotel RS, which combines both content and user similarity to calculate a relevance score for each review. In particular, we believe that better user similarity measures can be developed by taking into account ternary relationships such as those in our *co-staying network* [1]. Specifically, we are investigating connections between travelers who: 1) booked the same hotel(s); 2) stayed in similar hotels (e.g., that are part of the same chain); 3) have a well-defined set of topics that they mention frequently in their reviews; and 4) exhibit a similar rating behavior. Furthermore, we suggest a method for displaying a subset of personalized reviews graphically using Sankey diagrams. Allowing users to explore the multimode relationships could be considered as an additional form of explaining recommendations [20]. As future work, we aim to evaluate empirically whether these approaches, combined, increase users' understanding of the reasons behind recommending a specific hotel. We expect that such an outcome would, in turn, have a positive effect on the transparency and perceived trustworthiness of hotel RS.

Although not specifically discussed in this paper, methods for visualizing user opinions could be of interest also to hotel managers. In combination with interactive mechanisms, such as the ones suggested in the previous section, these graphical representations could provide a clearer picture of the feedback that guests typically write. This could help monitor and focus on areas that require improvement, i.e. topics with numerous negative mentions. The usefulness of these methods in other domains, such as data analytics or visualization RS [22], should also be investigated further.

ACKNOWLEDGMENTS

This work is supported by the German Research Foundation (DFG) under grant No. GRK 2167, Research Training Group "User-Centred Social Media".

REFERENCES

[1] Catalin-Mihai Barbu and Jürgen Ziegler. 2017. Co-Staying: A social network for increasing the trustworthiness of hotel recommendations. In *RecTour '17*. CEUR-WS Vol. 1906, 35–39.

[2] Guanliang Chen and Li Chen. 2015. Augmenting service recommender systems by incorporating contextual opinions from user reviews. *User Modeling and User-Adapted Interaction* 25, 3 (2015), 295–329.

[3] Cecilia di Sciascio, Vedran Sabol, and Eduardo E Veas. 2015. uRank: Exploring document recommendations through an interactive user-driven approach.. In *IntRS '15*. CEUR-WS Vol. 1438, 29–36.

[4] Tim Donkers, Benedikt Loepp, and Jürgen Ziegler. 2018. Explaining recommendations by means of user reviews. In *ExSS '18*. CEUR-WS Vol. 2068.

[5] Jan Feuerbach, Benedikt Loepp, Catalin-Mihai Barbu, and Jürgen Ziegler. 2017. Enhancing an Interactive Recommendation System with Review-based Information Filtering. In *IntRS '17*. CEUR-WS Vol. 1884, 2–9.

[6] F Maxwell Harper, Funing Xu, Harmanpreet Kaur, Kyle Condiff, Shuo Chang, and Loren Terveen. 2015. Putting users in control of their recommendations. In *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, 3–10.

[7] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9–27.

[8] Dmitry I Ignatov, Alexander Semenov, Daria Komissarova, and Dmitry V Gnatyshak. 2017. Multimodal Clustering for Community Detection. In *Formal Concept Analysis of Social Networks*. Springer, 59–96.

[9] Isaac Jones, Lei Tang, and Huan Liu. 2015. Community discovery in multi-mode networks. In *User Community Discovery*. Springer, 55–74.

[10] Bart P Knijnenburg, Martijn C Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 441–504.

[11] Joseph A Konstan and John Riedl. 2012. Recommender systems: from algorithms to user experience. *User Modeling and User-Adapted Interaction* 22, 1-2 (2012), 101–123.

[12] Johannes Kunkel, Benedikt Loepp, and Jürgen Ziegler. 2017. A 3D item space visualization for presenting and manipulating user preferences in collaborative filtering. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, 3–15.

[13] Bo Long, Zhongfei Mark Zhang, Xiaoyun Wu, and Philip S Yu. 2006. Spectral clustering for multi-type relational data. In *Proceedings of the 23rd International Conference on Machine Learning*. ACM, 585–592.

[14] Samaneh Moghaddam, Mohsen Jamali, and Martin Ester. 2011. Review recommendation: personalized prediction of the quality of online reviews. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*. ACM, 2249–2252.

[15] Sayooran Nagulendra and Julita Vassileva. 2014. Understanding and controlling the filter bubble through interactive visualization: a user study. In *Proceedings of the 25th ACM International Conference on Hypertext and Social Media*. ACM, 107–115.

[16] Denis Parra, Peter Brusilovsky, and Christoph Trattner. 2014. See what you want to see: visual user-driven approach for hybrid recommendation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*. ACM, 235–240.

[17] Christian Richthammer and Günther Pernul. 2016. Explorative analysis of recommendations through interactive visualization. In *International Conference on Electronic Commerce and Web Technologies*. Springer, 46–57.

[18] Rashmi Sinha and Kirsten Swearingen. 2002. The role of transparency in recommender systems. In *CHI'02 Extended Abstracts on Human Factors in Computing Systems*. ACM, 830–831.

[19] Tabea Tietz, Jörg Waitelonis, Joscha Jäger, and Harald Sack. 2014. Smart Media Navigator: Visualizing recommendations based on linked data. In *International Semantic Web Conference (Industry Track)*. CEUR-WS Vol. 1383.

[20] Nava Tintarev and Judith Masthoff. 2015. Explaining recommendations: Design and evaluation. In *Recommender Systems Handbook*. Springer, 353–382.

[21] Wenting Tu, David W Cheung, and Nikos Mamoulis. 2017. More focus on what you care about: Personalized top reviews set. *Neurocomputing* 254 (2017), 3–12.

[22] Manasi Vartak, Silu Huang, Tarique Siddiqui, Samuel Madden, and Aditya Parameswaran. 2017. Towards visualization recommendation systems. *ACM SIGMOD Record* 45, 4 (2017), 34–39.

[23] Katrien Verbert, Denis Parra, and Peter Brusilovsky. 2014. The effect of different set-based visualizations on user exploration of recommendations. In *IntRS '14*. CEUR-WS Vol. 1253, 37–44.

[24] Krist Wongsuphasawat and David Gotz. 2012. Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2659–2668.

[25] Mohammad Amin Yazdi, André Calero Valdez, Leonhard Lichtschlag, Martina Ziefle, and Jan Borchers. 2016. Visualizing opportunities of collaboration in large research organizations. In *International Conference on HCI in Business, Government and Organizations*. Springer, 350–361.

[26] Yongfeng Zhang. 2015. Incorporating phrase-level sentiment analysis on textual reviews for personalized recommendation. In *Proceedings of the 8th ACM International Conference on Web Search and Data Mining*. ACM, 435–440.

