

Increasing the Trustworthiness of Recommendations by Exploiting Social Media Sources

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

ABSTRACT

Current recommender systems mostly do not take into account as well as they might the wealth of information available in social media, thus preventing the user from obtaining a broad and reliable overview of different opinions and ratings on a product. Furthermore, there is a lack of user control over the recommendation process—which is mostly fully automated and does not allow the user to influence the sources and mechanisms by which recommendations are produced—as well as over the presentation of recommended items. Consequently, recommendations are often not transparent to the user, are considered to be less trustworthy, or do not meet the user's situational needs. This work will investigate the theoretical foundations for user-controllable, interactive methods of recommending, will develop techniques that exploit social media data in conjunction with other sources, and will validate the research empirically in the area of e-commerce product recommendations. The methods developed are intended to be applicable in a wide range of recommending and decision support scenarios.

Keywords

Recommender Systems; Interactive Recommending; User Interfaces; Trust.

1. INTRODUCTION AND MOTIVATION

For many years, recommender systems (RS) research has mostly focused on improving the algorithms used to generate recommendations, with aspects such as user control and the presentation of results receiving relatively little attention by comparison [9]. While the accuracy of RS has increased as a result of the introduction of more complex models and heuristics, users are finding it difficult to understand and influence the process through which recommendations are computed. The lack of transparency and minimal control have a detrimental effect on the perceived trustworthiness of such systems [5]. Against this background, further optimizing the already quite mature recommender algorithms alone is unlikely to yield significant additional benefits for the user. Instead, more user-oriented approaches to recommending are needed [9].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

RecSys '16, September 15 - 19, 2016, Boston, MA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-4035-9/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2959100.2959104>

Furthermore, current RS mostly do not take full advantage of the wealth of information that is available through social media channels. Collaborative filtering (CF) techniques have traditionally relied on user ratings to generate predictions. Other kinds of user-generated content that CF does not usually take into account, in particular textual contributions, tags, as well as user-user and user-item relationships, permeate the Web and can offer consumers a better overview of different opinions and ratings on a given product or service—potentially influencing the decisions that people make online. Since users tend to rely on the opinions of peers [4], it is reasonable to assume that they would also consider recommendations more trustworthy if the underlying RS were to augment them with information extracted from social media.

The aim of this thesis is to increase the trustworthiness of recommendations by developing more interactive approaches to recommending that enhance existing algorithms and methods with trust-related information extracted from social media sources. Furthermore, new methods for personalizing the presentation of recommended items with social media data will be produced to meet users' situational needs better. The research goals can be summarized as follows:

- Investigate and model users' decision making behavior with recommendations;
- Extract trust-related information from social media sources;
- Develop methods for providing users with trust-related information;
- Provide users with adaptive techniques for exploring trust information in social media and semantic data.

The remainder of this paper is structured as follows: Section 2 describes the state of the art with respect to the trustworthiness of recommendations, the personalization of the way in which these predictions are presented, and the relevant choice and decision support strategies that are being considered for RS. Section 3 details the proposed methods that will be applied in this research and in Section 4 the research plan is introduced. Section 5 concludes the paper by summarizing the expected contributions and next steps.

2. RELATED WORK

Trustworthiness of recommendations: There have been several studies on how humans build trust in RS. Herlocker et al. [7] found that providing explanations helps users understand the reasons that led the system to recommend certain items. By increasing transparency, the perceived trustworthiness of recommendations also increases. More recently, Pu et al. [13] proposed a trust model for RS that aims to raise the perceived trustworthiness of predictions by explaining the advantages and tradeoffs of each recommended item. Harman et al. [5] observed that people reacted more strongly to inaccurate predictions when these were

personalized: Their trust in the RS decreased more strongly and they were more dissatisfied with the results in comparison with the case where personalization was not used. Transparency and user control are often cited as having a significant effect on the perceived trust of RS and on the users' overall satisfaction with the technology (see, e.g., [9], [6], and [15]). This work will build upon these findings by investigating the effects that social media data has on recommendations.

Besides the factors that influence users' trust in RS, a separate line of research studies how trust in *other users* can be leveraged to improve the accuracy of CF techniques [10]. Such approaches typically rely on people marking other social network users as being trustworthy (or not). As it will be shown in the next section, this kind of information might be unavailable in the type of scenarios that this research proposes to address. In principle, however, it could be possible to extract trust measurements from social media interactions.

Personalizing the presentation of recommendations: Aspects such as what information to present, how to present it and when, and how much of it to present for any given prediction are important when discussing interactivity in RS. Prior work has, for instance, explored the effects of using different types of result lists and combinations of text with images on the persuasiveness of recommendations [12]. Other researchers suggested a model for timing recommendations [1] or determined the number of results that leads to high choice satisfaction without increasing choice difficulty [2]. However, most of this work does not focus specifically on interactive RS. Some interactive approaches to visualizing recommendations have been proposed, such as those implemented in TasteWeights [3] and TalkExplorer [17]. These systems afford a certain degree of control over the recommendation process to elicit feedback and preferences as well as to increase transparency. Personalizing the presentation of recommendations is typically achieved through explanations ([16], [14]). The proposed research will advance the state of the art in two ways. First, integrating social media data more tightly into the presentation of the recommended items will open up novel interaction possibilities for users. Second, the effect that such an enriched presentation has on the user's perceived trustworthiness of the recommendations will be evaluated empirically.

Choice and decision support for RS: Recent research suggests that novel approaches in RS can also stem from understanding how people make choices. Jameson et al. [8] performed a comprehensive analysis of this topic and proposed several ways in which recommender technology can be applied to support the human processes of choice and decision making. By giving users more control, they become involved in the decision-making process. Thus, recommendations also become a form of choice support. The authors note that typical support strategies for RS are "advise about processing" (i.e. recommending a particular course of action) and "evaluate on behalf of the chooser" (i.e. recommending an item or set of items). This work aims to investigate other choice support strategies that are not specifically related to recommendation technologies, such as "combine and compute" (i.e. performing complex calculations in order to present more relevant information) and "design the domain" (i.e. adapting the interface to make it easier for the user to choose between several options). To the author's knowledge, these aspects have not yet been explored in detail.

The specific research questions that will be addressed are:

- RQ1: What methods can be used to extract trust-related information from social media sources?

- RQ2: What measures for user similarity can be defined, taking into account factors such as social relationships, rating behavior, or elicited preferences?
- RQ3: How can the presentation of recommendations be adapted and personalized to increase users' trust?

3. PROPOSED METHODS

In order to explore the research questions, one specific domain was chosen in which the trustworthiness of recommendations is especially important, namely hotel booking. The reasons for selecting this domain are threefold: 1) There is a higher degree of risk associated with such choices (in comparison with, e.g., the frequently used domain of movie recommendations); 2) the items have a reasonable set of attributes that have to be taken into account; and 3) there is a large body of user-generated content, in the form of reviews, photos, tags, and ratings.

While analyzing the output of several online accommodation booking websites, it already became apparent that there is a lot of room for improving the way in which recommendations are presented to users. For instance, although most websites allow users to specify whether they are traveling for business or leisure, the effect of this choice on the presentation of individual recommendations is typically minimal. The argument put forth here is that the choices that people make during their search—destination, number of nights, desired amenities, purpose of travel, etc.—should influence the presentation of the output. Specifically, the presentation could be personalized not only in terms of the recommended items, but also tailored precisely to support each user's needs. To give an example, stating a preference for "fitness center" could lead to information such as opening hours, available machines, and pricing information being displayed more prominently.

When booking accommodation, each person has a number of expectations about what the ideal place ought to be like. Some are precise (e.g., the room must have soft beds), while others may be vaguer (e.g., the hotel should be close to public transport). Still others may apply only in specific situations, such as when a person travels for work (e.g., there must be a writing desk in the room). Expectations often evolve as the user becomes more experienced at traveling, or her social status changes over time. The proposed argument is that recommendations that satisfy users' expectations and are presented in a personalized fashion shall be considered more trustworthy. For a RS to achieve this, it must first be able to infer expectations from information available in user profiles, such as user ratings, reviews, tags (i.e. explicit feedback), interaction behavior (i.e. implicit feedback), and user relationships (from social media). Reviews, especially, can be a very rich source of information about what a user liked or disliked about a specific stay in a hotel. Such information can allow the system to learn more about the person who wrote them, in addition to its being a resource for other users. After returning from a trip, those who leave reviews will often highlight amenities that influenced—either positively or negatively—their stay in the hotel. (As a result, some online booking websites also made it possible for users to give separate ratings to frequently-mentioned hotel characteristics such as Wi-Fi quality.) Furthermore, reviews will often contain references to amenities that are not available as filters when selecting the hotel ("soft bed" being a good example).

By applying methods such as text mining and text analysis, an algorithm will be developed that extracts hotel attributes mentioned in reviews, calculates the user's degree of (dis)satisfaction based on her statements, and builds basic expectations that can be exploited by the RS during the generation and presentation of

recommendations (Figure 1). Implicit feedback, such as discarded recommendations, could also be exploited to identify “deal breakers”, namely attributes (or combinations thereof) that the user wants to avoid.

Extending the above procedure to an entire user base will allow RS to personalize and adapt the output by showing more relevant details about recommended items. To illustrate this point, consider someone who is interested in venues that offer good Wi-Fi connectivity. When browsing recommendations, she might find it useful to read reviews that specifically mention aspects such as connection speed and signal strength or those that give an overall quality assessment. To facilitate comparison, this information could, for example, be presented in the form of a graphical scale that shows the proportion of people who rated the internet connection positively versus those who rated it negatively. By taking into account expectations, a RS will be able to identify attributes that are most important to users and subsequently enhance the recommendations with such personalized summaries, thereby increasing their trustworthiness.

In addition to summarizing information related to user expectations, perhaps equally important is to make people aware of potential deal breakers they might not be aware of (Figure 2). If a significant proportion of guests have mentioned in their reviews that, e.g., the breakfast buffet was modest, this information could be passed along to the user in the form of advice [8]. Moreover, by allowing users to give feedback on these summaries, the RS could refine the user models and improve the presentation of the recommendations.

By mining past bookings and reviews, a complex network consisting of users, hotels, and hotel attributes can be created. This will allow to identify, with greater accuracy, which items a user is likely to find attractive based on the attributes mentioned in her reviews as well as in reviews of similar users [11]. At the same time, the system could extract and present, for each recommended item, the experiences of other people who are interested in the same combination of attributes as the current user. Such a network, which might be referred to as “co-staying in hotels”, will further facilitate the discovery of users with similar expectations and who have already stayed in the same hotel or in hotels with similar characteristics to the one that a person is considering. It is expected that this could thus introduce novel ways of interacting with RS. For example, someone who has a strong preference for soft beds will be able to explore the opinions of others who share this preference. Taking this a step further, the user might then ask herself, “What *other* preferences do such people have?” Personalizing the presentation to facilitate the exploration of similar people’s expectations could allow the user to discover new preferences that she would not have considered otherwise.

Giving users more control over the recommendation process and over the way in which these recommendations are presented also offers opportunities to experiment with learning by choosing, as proposed by Jameson et al. [8]. There is an attractive relationship between these two aspects. If a user already has fixed preferences, then user control means enabling her to put these preferences into practice (e.g., set the parameters for determining the relative weights of hotel attributes). But if she does not, it is a matter of enabling her to experiment to find out in the first place what she likes (as well as to put the results into practice subsequently). The types of appropriate support are different in the two cases; however, the second case is less well understood. This research will provide




Figure 1 – Extracting expectations from reviews

the basis for studying novel interaction techniques that help users learn while making choices about the recommendations.

An initial model showing the interactions between the three components that have been discussed throughout this section, namely user profiles, hotel attributes, and hotels, is depicted in Figure 3: Information stored in a user’s profile is exploited to identify individual expectations and similar users. These, in turn, help the system to select the appropriate hotel attributes. Hotel recommendations generated from these attributes are refined and reordered based on the reviews produced by similar users. Finally, the personalized recommendations are presented to the user.

4. RESEARCH PLAN

In order to implement, refine, and evaluate the methods proposed in the previous section, an adequate dataset is needed. For the hotel booking scenario, this translates into: 1) A sufficiently large number of hotels, ideally covering several major European cities; and 2) a diverse set of reviews contributed by people travelling for different purposes. The first requirement ensures a representative and interesting subset of hotels, covering as many types of amenities as possible. The second requirement is intended to maximize the variety of user opinions and arguments, thereby increasing the range of possible user expectations. Since most datasets available online do not contain user-generated data such as reviews, ratings, and tags, it was necessary to build a custom crawler for collecting this information into a database. The online booking website Booking.com was chosen as the data source because it only allows verified hotel guests to submit reviews. The collected dataset comprises 576,506 reviews, written by 219,699 users who have stayed in at least one of the 10,154 establishments in 5 major European cities¹. This should provide an adequate testbed for developing algorithms to mine reviews and user profiles necessary for extracting trust-related information (cf. RQ1).

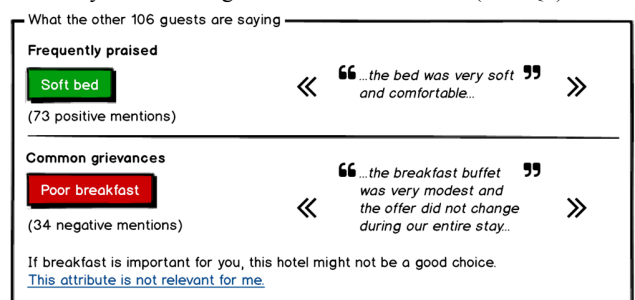


Figure 2 - Making the user aware of potential problems

Given that similar users play an important role in the model introduced in the previous section, developing accurate algorithms for identifying and ranking them is critical for success of the proposed approach. Existing metrics for computing user similarity, which typically consider rating behaviors, will be extended by taking into account user expectations, attributes mentioned in reviews, and co-staying (cf. RQ2).

In parallel to building the knowledge base, use cases for the hotel booking scenario are being prepared, which will serve as the basis

¹ Barcelona, Berlin, Brussels, London, and Rome

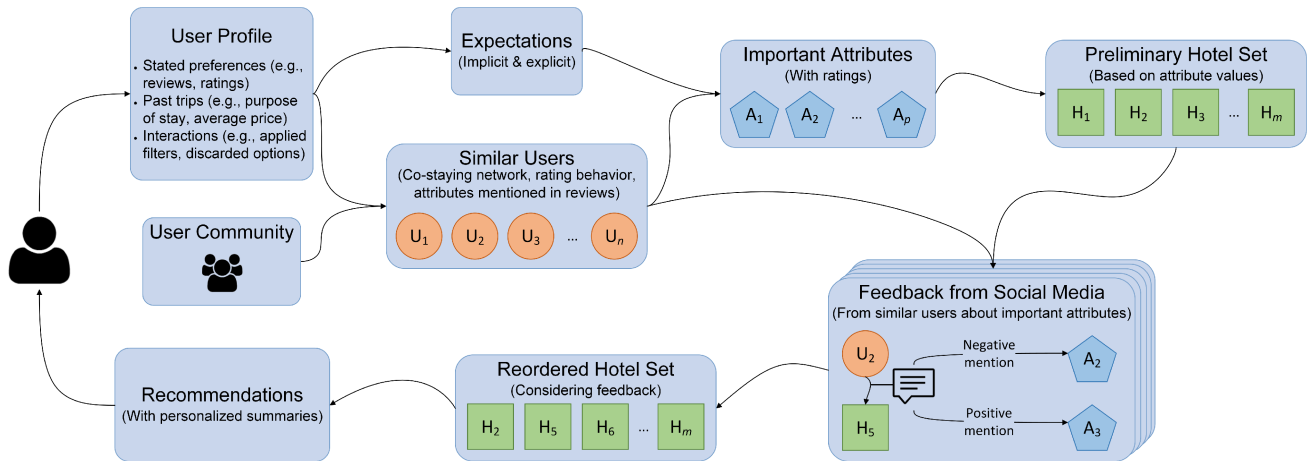


Figure 3 - Initial model for enhancing recommendations by exploiting social media data

for the initial user interface mockups. The purpose of the mockups is to decide what adapted information will be presented, what interaction possibilities (i.e. to related information) will be available, and the interactive mechanisms through which they will be achieved. Once finished, the mockups will be evaluated in a user study to identify usability problems and to further refine the above-mentioned model. Subsequent implementations and iterative testing of the recommendation technology will investigate the theoretical foundations for user-controllable, interactive methods of recommending, will develop techniques to exploit social media data to increase the trustworthiness of recommendations, and will validate the research empirically (cf. RQ3).

5. CONCLUDING REMARKS

This paper proposes several methods through which current RS could be improved further by 1) leveraging the growing corpus of user-generated content to extract relevant information, and 2) personalizing the presentation of the recommended items so that they are perceived by users as more trustworthy. The next steps will focus on methods for extracting relevant information from the mined reviews and on improving existing user similarity measures.

6. ACKNOWLEDGEMENTS

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

7. REFERENCES

- [1] Betzalel, N. D., Shapira, B., & Rokach, L. "Please, not now!" A model for timing recommendations. In *Proc. RecSys '15*, ACM (2015), 297–300.
- [2] Bollen, D., Knijnenburg, B. P., Willemsen, M. C., & Graus, M. Understanding choice overload in recommender systems. In *Proc. RecSys '10*, ACM (2010), 63–70.
- [3] Bostandjiev, S., O'Donovan, J., & Höllerer, T. Tasteweights: a visual interactive hybrid recommender system. In *Proc. RecSys '12*, ACM (2012), 35–42.
- [4] Flanagan, A. J., & Metzger, M. J. Trusting expert-versus user-generated ratings online: The role of information volume, valence, and consumer characteristics. *Computers in Human Behavior*, 29(4), 1626–1634.
- [5] Harman, J. L., O'Donovan, J., Abdelzaher, T., & Gonzalez, C. Dynamics of human trust in recommender systems. In *Proc. RecSys '14*, ACM (2014), 305–308.
- [6] Harper, F. M., Xu, F., Kaur, H., Condiff, K., Chang, S., & Terveen, L. Putting users in control of their recommendations. In *Proc. RecSys '15*, ACM (2015), 3–10.
- [7] Herlocker, J. L., Konstan, J. A., & Riedl, J. Explaining collaborative filtering recommendations. In *Proc. CSCW '00*, ACM (2000), 241–250.
- [8] Jameson, A., Willemsen, M. C., Felfernig, A., Gemmis, M. de Lops, P., Semeraro, G., & Chen, L. Human Decision Making and Recommender Systems. *Recommender Systems Handbook*, Springer US (2015), 611–648.
- [9] Konstan, J. A., & Riedl, J. Recommender systems: From algorithms to user experience. *User Mod. and User-Adapted Interaction* 22, 1-2 (2012), 101–123.
- [10] Massa, P., & Avesani, P. 2007. Trust-aware recommender systems. In *Proc. RecSys '07*. ACM (2007), 17–24.
- [11] McAuley, J., & Leskovec, J. Hidden factors and hidden topics: Understanding rating dimensions with review text. In *Proc. RecSys '13*, ACM (2013), 165–172.
- [12] Nanou, T., Lekakos, G., & Fouskas, K. The effects of recommendations' presentation on persuasion and satisfaction in a movie recommender system. *Multimedia Systems* 16, 4-5 (2010), 219–230.
- [13] Pu, P., Chen, L., & Hu, R. Evaluating recommender systems from the users' perspective: Survey of the state of the art. *User Mod. and User-Adapted Interaction* 22, 4-5 (2012), 317–355.
- [14] Pu, P., & Chen, L. Trust-inspiring explanation interfaces for recommender systems. *Knowledge-Based Systems*, 20(6) (2007), 542–556.
- [15] Sinha, R., & Swearingen, K. The Role of Transparency in Recommender Systems. In *Proc. CHI EA '02*, ACM (2002), 830–831.
- [16] Tintarev, N., & Masthoff, J. Designing and evaluating explanations for recommender systems. *Recommender Systems Handbook*. Springer US (2011), 479–510.
- [17] Verbert, K., Parra, D., Brusilovsky, P., & Duval, E. Visualizing recommendations to support exploration, transparency and controllability. In *Proc. IUI '13*, ACM (2013), 351–362.