

SocRecM: A Scalable Social Recommender Engine for Online Marketplaces

Emanuel Lacic
Graz University of Technology
Graz, Austria
elacic@know-center.at

Dominik Kowald
Know-Center
Graz, Austria
dkowald@know-center.at

Christoph Trattner
Know-Center
Graz, Austria
ctrattner@know-center.at

ABSTRACT

This paper presents work-in-progress on *SocRecM*, a novel social recommendation framework for online marketplaces. We demonstrate that *SocRecM* is easy to integrate with existing Web technologies through a RESTful, scalable and easy-to-extend service-based architecture. Moreover, we reveal the extent to which various social features and recommendation approaches are useful in an online social marketplace environment.

Categories and Subject Descriptors

H.2.8 [Database Management]: Database Applications—*Data mining*; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—*Information filtering*

Keywords

social recommender engine; online marketplaces; recommender framework; recommender evaluation; Apache Solr

1. INTRODUCTION

Recommender systems aim at helping users find relevant information in an overloaded information space. Although various recommender frameworks are available nowadays, there is still a lack of frameworks that address such important aspects in recommender systems research as: easy integration into an existing infrastructure, scalability, hybridization and social data integration. To tackle these issues, we implemented *SocRecM*, a scalable and easy-to-integrate online social recommender framework whose purpose is not only to decrease the workload of developers via an easy-to-use framework but also to provide recommendation algorithms that utilize social data obtained from various data sources (e.g., Facebook, Twitter, etc.).

2. SYSTEM OVERVIEW

The first prototype of *SocRecM* was implemented in Java and can be found online as open-source software¹. Figure 1 shows the

¹<https://github.com/learning-layers/SocRec>

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s). *HT'14*, September 1–4, 2014, Santiago, Chile. ACM 978-1-4503-2954-5/14/09. <http://dx.doi.org/10.1145/2631775.2631783>.

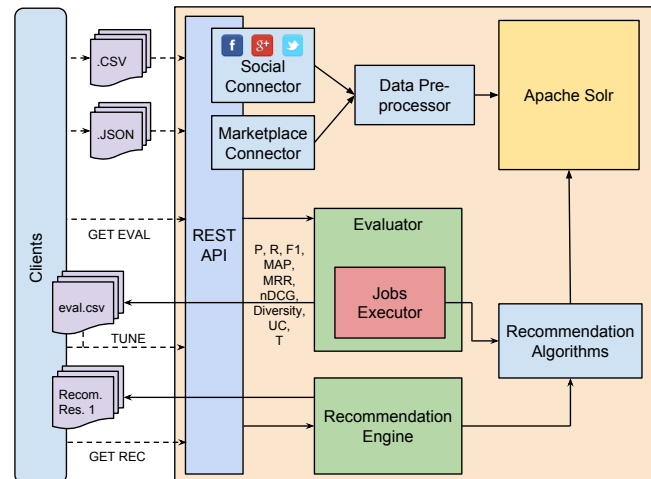


Figure 1: *SocRecM* system architecture.

system architecture of *SocRecM*. As featured, the engine can easily be integrated into a RESTful API for communicating with client applications. The *SocRecM* API provides methods for uploading marketplace and social data into the engine and for querying resource recommendations and benchmarking results. The marketplace and social data are gathered via respective connectors and pre-processed to be indexed by Apache Solr, which in turn offers powerful search and content analyzing functionalities (e.g., facets or MoreLikeThis queries) that are used by the recommendation algorithm implementations. Currently, *SocRecM* contains four types of algorithms to recommend resources (in our case products) to users: MostPopular (MP), Collaborative Filtering (CF) [6], Content-based (C) [5] and Hybrid Recommendations (CCF) [1] (see also [4]). These algorithms are calculated based on either marketplace features, such as purchases (CF_p), title (C_t) and description (C_d) or social features, such as likes (CF_l), comments (CF_c), interactions (CF_{in}), social stream content data (C_{st}), groups (CF_g) and interests (CF_i). Additionally, the algorithms and data features are incorporated into hybrid algorithms (CCF_m for marketplace and CCF_s for social features).

The recommendation algorithms are invoked by the evaluator component and its jobs executor in order to evaluate them with respect to various IR metrics (e.g., Recall (R), Precision (P), nDCG, User Coverage (UC), etc.) [2, 7]. The evaluation results are further used in *SocRecM* to tune the parameters of the algorithms, especially in the case of the hybrid approaches. Furthermore, the algo-

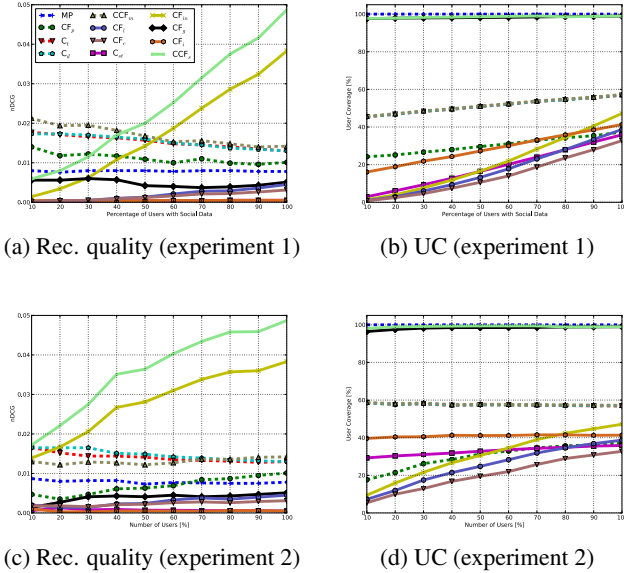


Figure 2: Results of the two experiments in terms of recommendation accuracy (nDCG) and User Coverage (UC).

gorithms are called by the recommendation engine in order to forward the recommended resources back to the client.

3. RESULTS

To demonstrate the effectiveness of our approach under different scenarios we conducted two “virtual” experiments with a social dataset gathered from the virtual world of SecondLife, which provides both: detailed marketplace purchase data and social data collected from users in the virtual world (similar as in [9]). The dataset contains 126,356 users, 122,360 products, 265,274 purchases, 1,839,783 social interactions, 510,145 social stream contents, 260,137 groups and 88,371 interests and was split into a training and test set using the method proposed in [4]. The experiments described below were conducted using an IBM System x3550 server with two 2.0 GHz six-core Intel Xeon E5-2620 processors, a 1TB ServeRAID M1115 SCSI Disk and 128 GB of RAM using one Apache Solr 4.3.1 instance in the back-end.

The first experiment demonstrated how the recommendation quality in terms of the nDCG and User Coverage (UC) metrics [3] could be improved if the users’ social data were provided in addition to the marketplace data. To simulate this, we extracted all users from our dataset that have both the marketplace and social data (10,996 users). Out of this subset, we randomly selected 10% of users and replaced them with users that only provided marketplace data. We continued until 100% of the data consisted of users with only marketplace data in their profiles. As demonstrated in plots (a) and (b) of Figure 2, the recommendation approaches based on marketplace data alone are fairly constant regardless of how many users have social data in their profiles. This was expected since the number of users neither increased nor decreased throughout the experiment. However, it is apparent that although the recommenders based on the users’ interactions perform best (CCF_s and CF_{in}), they depend greatly on the number of social profiles in the dataset. An interesting finding in this context was that the content-based approach based on the users’ social stream contents (CF_{st}) performed as poorly as the recommender based on the users’ interests (CF_i).




Figure 3: Mean runtime of the algorithms in SocRecM.

In the second experiment we simulated a cold-start scenario for a new social marketplace system, under which we assumed that all new users provided social data as it was the case with the start of Spotify, to determine if and how the recommendation quality would be affected by an increasing number of social users. In order to conduct this experiment, we once again extracted all users from the dataset with both marketplace and social data (10,996) and eliminated the rest. Using these data, we created 10 different sets with an increasing number of users and evaluated the recommendation approaches on them, as shown in plots (c) and (d) in Figure 2. As the number of users with marketplace and social data increased, the approaches based on the social data (CCF_s and CF_{in}) delivered a much higher prediction quality than those that used only the marketplace data (e.g., CF_p). Moreover, the user coverage of the hybrid approaches based on the social data (CCF_s) was much higher than the one based on the marketplace data (CCF_m). Comparing CF_p with CF_{in}, shows that although both algorithms have a higher user coverage with more users, the social features ultimately provide better results.

In addition to the recommendation quality, we compared the mean runtime (i.e., the time needed to calculate recommendations for a user) of the recommendation approaches shown in Figure 3. In general these results demonstrate that SocRecM is capable of providing near real-time recommendations for users since the maximum mean test time was only 58 milliseconds for the hybrid approaches CCF_m and CCF_s.

4. FUTURE WORK

In the future, we plan to further extend our framework using different approaches to make sense of social data provided by users in a social marketplace environment. For example, we would like to implement a topic modeling approach based on LDA to combine comments with purchase data or to derive topics from the user’s social stream to calculate similarities between users. Furthermore, we are interested in extending our framework to generate recommendations based on the user’s geo-location data (e.g., [8]). Last but not least, we are interested in developing novel hybridization approaches for diverse social data sources to further increase the predictive power of our recommender framework.

Acknowledgments: This work is supported by the Know-Center and the EU funded project Learning Layers (Grant Nr. 318209). The Know-Center is funded within the Austrian COMET Program - Competence Centers for Excellent Technologies - under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency (FFG).

5. REFERENCES

- [1] R. Burke. Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction*, 12(4):331–370, 2002.
- [2] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.
- [3] K. Järvelin and J. Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Trans. Inf. Syst.*, 20(4):422–446, Oct. 2002.
- [4] E. Lacic, D. Kowald, D. Parra, M. Kahr, and C. Trattner. Towards a scalable social recommender engine for online marketplaces: The case of apache solr. In *Proc., WWW '14*. ACM, 2014.
- [5] M. J. Pazzani and D. Billsus. Content-based recommendation systems. In *The adaptive web*, pages 325–341. Springer, 2007.
- [6] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen. Collaborative filtering recommender systems. In *The adaptive web*, pages 291–324. Springer, 2007.
- [7] B. Smyth and P. McClave. Similarity vs. diversity. In D. Aha and I. Watson, editors, *Case-Based Reasoning Research and Development*, volume 2080 of *LNCS*, pages 347–361. Springer, 2001.
- [8] M. Steurer and C. Trattner. Acquaintance or Partner? Predicting Partnership in Online and Location-Based Social Networks. In *Proc. ASONAM'13*. IEEE/ACM, 2013.
- [9] Y. Zhang and M. Pennacchiotti. Predicting purchase behaviors from social media. In *Proceedings of the 22Nd International Conference on World Wide Web, WWW '13*, pages 1521–1532, 2013.