

AFEL-REC: A Recommender System for Providing Learning Resource Recommendations in Social Learning Environments

Dominik Kowald, Emanuel Lacic, Dieter Theiler & Elisabeth Lex

Know-Center GmbH & Graz University of Technology, Austria
{dkowald,elacic,dtheiler}@know-center.at,elisabeth.lex@tugraz.at

ABSTRACT

In this paper, we present preliminary results of AFEL-REC, a recommender system for social learning environments. AFEL-REC is build upon a scalable software architecture to provide recommendations of learning resources in near real-time. Furthermore, AFEL-REC can cope with any kind of data that is present in social learning environments such as resource metadata, user interactions or social tags. We provide a preliminary evaluation of three recommendation use cases implemented in AFEL-REC and we find that utilizing social data in form of tags is helpful for not only improving recommendation accuracy but also coverage. This paper should be valuable for both researchers and practitioners interested in providing resource recommendations in social learning environments.

KEYWORDS

Social Recommender Systems; Social Learning Environments; Analytics for Everyday Learning; Collaborative Filtering; Coverage

1 INTRODUCTION

Recommender systems aim to predict if a specific user will like a specific resource. To do so, recommender systems analyze past usage behavior (e.g., clicks or likes) with the goal to generate a personalized list of potentially relevant resources [15]. Nowadays, recommender systems are part of many applications, such as online marketplaces (e.g., Amazon), movie streaming services (e.g., Netflix), job portals (e.g., LinkedIn), and Technology Enhanced Learning (TEL) environments (e.g., Coursera).

Especially in the field of TEL, recommender systems have become an important research area over the past decade [1]. One of the many examples in this area is the CoFind system [2], which guides learners to resources that were found useful by other learners in the past. Other examples include the work of [13], who proposed a recommendation approach with query extraction mechanisms or the work of [4], who enhanced Collaborative Filtering (CF) [5] by taking into account the learner's evolution in time. Another recent strand is the research on context-aware recommender approaches for TEL. Here, contextual information, such as the location [19], are incorporated into the recommendation process. In this respect, social learning environments [18], which aim to support users in learning through the observation of other users' behaviors, bear great potential for recommender systems as they provide a vast amount of social information. This includes, for example, friendship connections, group memberships or social tags, which are

freely-chosen keywords used for collaboratively annotating learning resources [8]. Although a lot of recommender systems and algorithms are available in the TEL area, there is still the lack of research on recommender systems specifically tailored for social learning environments (such as e.g., [3]).

Therefore, in the course of the H2020 project Analytics for Everyday Learning (AFEL)¹, we have developed AFEL-REC, which is a recommender system for social learning environments. AFEL-REC is build upon a scalable software architecture to support various use cases for providing recommendations of learning resources in near real-time (see Section 2). We conducted a preliminary evaluation of AFEL-REC using data gathered from the Spanish social learning environment Didactalia² (see Section 3). Taken together, our contributions are twofold (see also Section 4):

- (1) We present AFEL-REC and use cases for providing resource recommendations in social learning environments.
- (2) We demonstrate that social information, such as social tags, can be used to improve the accuracy and coverage of recommendations in social learning environments.

We believe that our work contributes to the under-researched portfolio of recommender systems for social learning environments. Furthermore, we present an overview of use cases and preliminary evaluation results, which should be valuable for both researchers and practitioners interested in providing resource recommendations in social learning environments.

2 APPROACH

In this section, we present our AFEL-REC system by providing a detailed description of its software architecture as well as potential use cases that can be realized for recommending resources in social learning environments.

2.1 System Overview

The software architecture of AFEL-REC is based on the scalable recommendation framework ScaR³ [11]. It is depicted in Figure 1 and consists of the following main modules:

Service Provider (SP). The SP acts as a proxy for social learning environments to access AFEL-REC. Thus, it provides REST-based Web services to enable clients to query recommendations and to add new data (e.g., user interactions or learning resources) to the recommender system.

SIR'18@CIKM, October 2018, Turin, Italy

© 2018 Copyright for this paper held by its authors. Copying permitted for private and academic purposes.

¹<http://afel-project.eu/>

²<https://didactalia.net/>

³<http://scar.know-center.tugraz.at/>

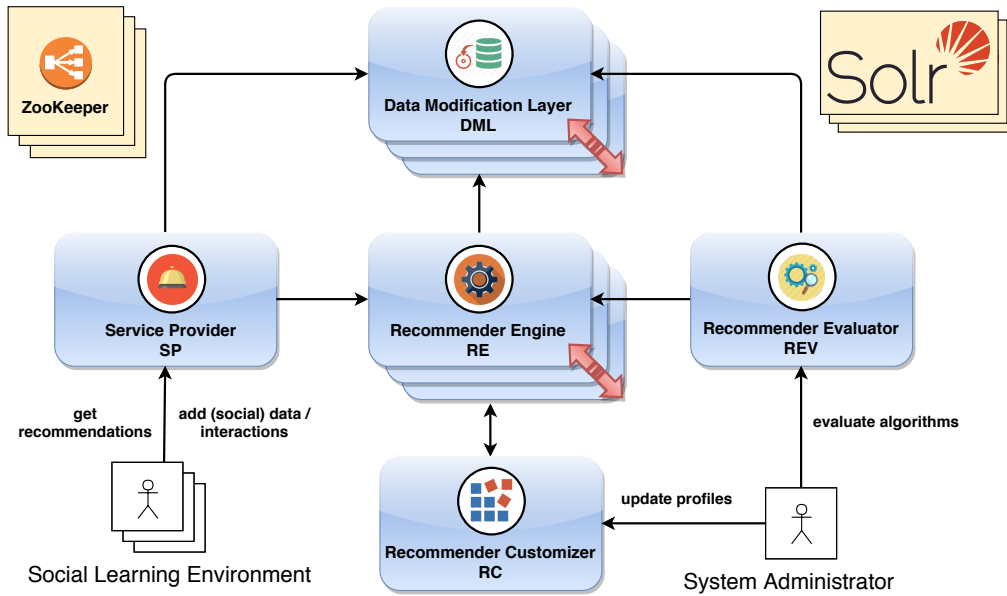


Figure 1: The scalable AFEL-REC software architecture based on the powerful open-source frameworks Apache Solr and Apache ZooKeeper. The main modules of AFEL-REC communicate via REST-based Web services.

Data Modification Layer (DML) & Solr. The DML encapsulates all CRUD operations (i.e., create, retrieve, update, delete) in one module and therefore, enables easy access to the underlying data backend. As depicted in Figure 1, AFEL-REC uses the high-performance search platform Apache Solr⁴. This data backend solution not only guarantees scalability and near real-time recommendations but also the support of multiple data sources. While most recommender systems rely on rating-based data, AFEL-REC is also capable of processing relevant social information such as tags.

Recommender Engine (RE). The RE is the heart of AFEL-REC and is responsible for calculating recommendations. As we are using Apache Solr, we can benefit from its build-in data structures for efficiently calculating user and resource similarities, and ranking resources based on their relevance for a specific recommendation context. In Section 2.2, we discuss possible use cases and algorithms that can be realized using AFEL-REC’s recommender engine.

Recommender Customizer (RC). The RC is used to change the parameters (e.g., the neighborhood size n) of the recommendation approaches on the fly. Thus, it holds a so-called recommendation profile for each approach, which can be accessed and changed by the system administrator. These changes are then broadcast to the RE to be aware of how a specific approach should be executed.

Recommender Evaluator (REV). The REV is responsible for evaluating the recommendation algorithms implemented in the RE. Thus, it can be executed to perform an offline evaluation with training/test set splits (see Section 3) or an online evaluation with A/B-tests.

ZooKeeper. We are using Apache ZooKeeper⁵ for handling the communication between the modules and for offering horizontal scaling. Thus, in cases in which we observe a high request load, we can start multiple instances of the same module (indicated by the arrows in the DML and RE modules).

2.2 Use Cases

Using this software architecture, AFEL-REC is capable of supporting seven use cases for providing recommendations in social learning environments:

UC1: Recommendation of Popular Resources in the Social Learning Environment. The first use case is a non-personalized one and is especially useful for new users of a social learning environment without any user interactions so far (i.e., cold-start users [16]). Thus, it is typically realized using a MostPopular algorithm. This approach recommends learning resources, which are weighted and ranked by the number of interactions. As mentioned, the Most-Popular approach is non-personalized and thus, each user will receive the same recommendations.

UC2: Recommendation of Resources That Like-Minded Users Have Interacted With. The second use case is a personalized one and thus, analyzes past user interactions to specifically tailor recommendations towards learners. Collaborative Filtering (CF) algorithms are typically chosen to realize such a use case [5]. CF approaches analyze the interactions between users and items (e.g., learning resources) and recommend those items to a given user that similar users have interacted with in the past. More specifically, in CF methods two users are treated as similar if they have liked the same items in the past. This in turn allows us to assume that these two users will also like the same (or similar) items in the future.

⁴<http://lucene.apache.org/solr/>

⁵<https://zookeeper.apache.org/>

UC3: Recommendation of Resources Based on Social Information. This use case is similar to UC2 but this time two users are treated as similar if they have shared some social information in the past. This social information can be friendship connections, group memberships or social tags. We hypothesize that social information is capable of providing a richer semantic representation of a user's interests than pure interaction data. Therefore, we also think that this should positively influence the coverage of the recommendations as validated in Section 3 of this paper.

UC4: Recommendation of Resources That are Similar to the Resources the User Has Interacted With. One disadvantage of UC2 and UC3 is that it can only be applied to resources, which already have user interactions or social information attached to them. This means that cold-start resources without any user interactions or social information cannot be recommended. To overcome this, UC3 aims at utilizing resource similarities for personalized recommendations by using a Content-based Filtering (CBF) approach [12]. CBF methods use resource features such as categories or description texts for calculating similarities between resources. Then, the most similar resources of the resources the given user has interacted with are recommended.

UC5: Recommendation of (Alternative) Resources for a Specific Resource. This use case is related to UC4 but provides contextualized recommendations instead of personalized ones. This means that recommendations are not based on the learner but based on a specific resource by finding alternative ones. Thus, similar to UC4, the most similar resources for the given resource are recommended using CBF.

UC6: Recommendation of (Alternative) Resources for a Specific User and a Specific Resource. The next use case also focuses on recommending alternative resources for a specific resource but this time in a personalized manner. Such use case can be implemented using a contextualized CF approach. This means that we search for similar users of the target user, who have also interacted with the target resource. Thus, this use case is similar to Amazon's "Users who bought this, also bought that" recommender.

UC7: Recommendation of Resources for a Specific User and a Specific Learning Goal. Finally, the last use case tackles recommendations in a personalized and adaptive manner. While UC1 to UC6 have focused on providing relevant recommendations, they have neglected the learning goal of the user. Such a learning goal could be the aim to focus on a specific topic or to receive more difficult learning resources. To realize adaptive recommendations, the suggested resources could be re-ranked using a feature boosting technique. For example, if the learning goal is to study more difficult resources, then resources with a higher complexity (e.g., measured via a readability score) should be boosted and easy ones should be down-graded in the recommendation list.

3 PRELIMINARY EVALUATION

In this section, we present preliminary evaluation results for AFEL-REC. The aim of this evaluation is twofold: (i) we want to show that AFEL-REC is capable of providing recommendations using data gathered from a real-world social learning environment, and (ii),

Number of interactions (i.e., clicks)	1,879,761
Number of users	1,274,858
Number of learning resources	35,346
Number of social tags	485,295
Average number of interactions per user	1.47
Average number of interactions per learning resource	53.18
Average number of tags per learning resource	13.73

Table 1: Statistics of our dataset, which was collected from the social learning environment Didactalia.

we want show that social information can enhance the prediction accuracy and coverage of recommendations. To do so, we focus on UC1 to UC3 presented in the previous section.

3.1 Data

We collected our data from the social learning environment Didactalia between the 26th of February 2017 until the 28th of May 2018. This included 1,879,761 user interactions (i.e., clicks on learning resources) by 1,274,858 users on 35,346 learning resources. This resulted in 1.47 interactions per user and 53.18 interactions per learning resource on average. Additionally, 485,295 social tags were applied to these learning resources, which resulted in 13.73 tags per resource on average. The full statistics of our dataset are shown in Table 1. To date, the only social information we are using in our data are tags but we are planning to extend this by also incorporating connections between users or group memberships [10].

3.2 Evaluation Method and Metrics

For evaluating AFEL-REC, we split our dataset into training and test sets. Therefore, we followed common practice in the research area of recommender systems and information retrieval by using the most recent 20% of interactions of each user for testing and the remaining 80% for training. This dataset splitting technique ensures that the chronological order of the data is preserved and thus, that the future is predicted based on past user interactions.

For measuring the accuracy of the recommendations, we use a rich set of metrics, namely Recall ($R@20$, measured for $k = 20$ recommended resources), Precision ($P@1$, for $k = 1$), F1-score ($F1@10$ for $k = 10$), Mean Reciprocal Rank ($MRR@20$, for $k = 20$), Mean Average Precision ($MAP@20$, for $k = 20$) and normalized Discounted Cumulative Gain ($nDCG@20$, for $k = 20$) [17]. Furthermore, we also report the coverage (C) of the recommendations, measuring the fraction of users for whom the algorithm is capable of producing any recommendations.

3.3 Recommendation Approaches

We evaluated UC1 - UC3 presented in Section 2.2 to show (i) the general usefulness of AFEL-REC for providing recommendations in social learning environments, and (ii) that social information in the form of tags is helpful for improving the recommendation accuracy and coverage. In the future, we will also evaluate UC4 - UC7.

The MP (MostPopular) approach refers to UC1 and is a non-personalized algorithm, which recommends the most frequently used learning resources in the system. This algorithm also works for cold-start users [16] and thus, should reach a UC of 100%.

Approach	R@20	P@1	F1@10	MRR@20	MAP@20	nDCG@20	C
UC1: MP	.007	.002	.002	.002	.002	.003	100%
UC2: CF _i	.046	.022	.012	.025	.026	.032	40%
UC3: CF _t	.070	.027	.016	.034	.035	.044	53%

Table 2: Preliminary results of our evaluation of AFEL-REC. We see that CF_t provides not only a better recommendation accuracy but also coverage (C) than CF_i, which means that social information in the form of tags is helpful for improving recommendations in social learning environments.

The CF_i algorithm refers to UC2 and calculates the neighborhood of a user on the basis of interaction data (i.e., clicks). Thus, two users are treated as similar if they have interacted with the same learning resources in the past [5]. Based on [6], we used a neighborhood size of $n = 20$ users.

Similar to CF_i, CF_t is a Collaborative Filtering-based approach but, as discussed in UC3, this one calculates the user neighborhood based on social tags. Thus, two users are treated as similar if they have used the same social tags in the past [14].

3.4 Results

The preliminary results of our evaluation are shown in Table 2. While the lowest accuracy estimates are provided by the unpersonalized MP approach (UC1), this approach also provides the highest UC with 100%. This means that MP is capable of providing recommendations for all users in the social learning environment.

When looking at the next algorithm, CF_i (i.e., CF on the basis of user interactions, UC2), we notice that this approach provides approx. 10 times higher accuracy values than MP. We contribute this to the personalization factor of the algorithm. However, one drawback of CF_i is the rather small UC of 40%, which means that it cannot generate any recommendations for 60% of the users.

Finally, CF_t (i.e., CF on the basis of social tags, UC3), provides not only the best results with respect to recommendation accuracy but also a larger coverage (C) than CF_i. This result shows that social information can indeed help to improve recommendations in social learning environments.

4 CONCLUSION AND FUTURE WORK

In this paper, we presented AFEL-REC, a recommender system for social learning environments. AFEL-REC is build upon a scalable software architecture based on powerful open-source frameworks such as Apache Solr or Apache ZooKeeper and thus, is capable of providing near real-time recommendations of learning resources. We have demonstrated the usefulness of AFEL-REC by discussing use cases that can be realized with it and by providing preliminary evaluation results using data gathered from the Spanish social learning environment Didactalia. Our evaluation results show that social information in the form of tags can be used to enhance the accuracy and coverage of learning resource recommendations. These findings should be of interest for both researchers and practitioners interested in providing resource recommendations in social learning environments.

Future Work. One limitation of this work is that we have only investigated tags as potential social information for learning resource recommendations. Thus, for future work, we plan to extend

AFEL-REC by also considering additional types of social information such as friendship connections or group memberships. Also, we plan to include our cognitive-inspired tag recommender algorithms, which have been specifically useful in the context of TEL [7, 9]. Finally, we will also evaluate the remaining five use cases, which we have not tackled so far. This evaluation will then be also conducted in an online manner using a project-wide user study.

Acknowledgments. The authors would like to thank Didactalia and the AFEL consortium. This work was supported by the Knowledge Center GmbH Graz (Austrian FFG COMET Program) and the EU-funded H2020 project AFEL (GA: 687916).

REFERENCES

- [1] Hendrik Drachler, Katrien Verbert, Olga C Santos, and Nikos Manouselis. 2015. Panorama of recommender systems to support learning. In *Recommender systems handbook*. Springer, 421–451.
- [2] Jon Dron, Richard Mitchell, Phil Siviter, and Chris Boyne. 2000. CoFIND—an experiment in n-dimensional collaborative filtering. *Journal of Network and Computer Applications* 23, 2 (2000), 131–142.
- [3] Sandy El Helou, Christophe Salzmann, and Denis Gillet. 2010. The 3A Personalized, Contextual and Relation-based Recommender System. *J. UCS* 16 (2010).
- [4] Mercedes Gomez-Albarran and Guillermo Jimenez-Diaz. 2009. Recommendation and students’ authoring in repositories of learning objects: A case-based reasoning approach. *International Journal of Emerging Technologies in Learning (iJET)* 4, 2009 (2009), 35–40.
- [5] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004), 5–53.
- [6] Simone Kopeinik, Dominik Kowald, Ilire Hasani-Mavriqi, and Elisabeth Lex. 2017. Improving Collaborative Filtering Using a Cognitive Model of Human Category Learning. *The Journal of Web Science* 2, 1 (2017).
- [7] Simone Kopeinik, Elisabeth Lex, Paul Seitlinger, Dietrich Albert, and Tobias Ley. 2017. Supporting Collaborative Learning with Tag Recommendations: A Real-world Study in an Inquiry-based Classroom Project. In *Proceedings of the Seventh International Learning Analytics and Knowledge Conference (LAK '17)*. ACM, New York, NY, USA, 409–418. DOI: <http://dx.doi.org/10.1145/3027385.3027421>
- [8] Dominik Kowald, Simone Kopeinik, Paul Seitlinger, Tobias Ley, Dietrich Albert, and Christoph Trattner. 2015. Refining frequency-based tag reuse predictions by means of time and semantic context. In *Mining, Modeling, and Recommending 'Things' in Social Media*. Springer, 55–74.
- [9] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal Effects on Hashtag Reuse in Twitter: A Cognitive-Inspired Hashtag Recommendation Approach. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. 1401–1410.
- [10] Emanuel Lacic, Dominik Kowald, Lukas Eberhard, Christoph Trattner, Denis Parra, and Leandro Balby Marinho. 2015. Utilizing online social network and location-based data to recommend products and categories in online marketplaces. In *Mining, Modeling, and Recommending 'Things' in Social Media*. Springer.
- [11] Emanuel Lacic, Matthias Traub, Dominik Kowald, and Elisabeth Lex. 2015. ScaR: Towards a Real-Time Recommender Framework Following the Microservices Architecture. In *Proceedings of LRSR2015 Workshop at RecSys 2015*.
- [12] Pasquale Lops, Marco De Gemmis, and Giovanni Semeraro. 2011. Content-based recommender systems: State of the art and trends. In *Recommender systems handbook*. Springer, 73–105.
- [13] Eleni Mangina and John Kilbride. 2008. Evaluation of keyphrase extraction algorithm and tiling process for a document/resource recommender within e-learning environments. *Computers & Education* 50, 3 (2008), 807–820.
- [14] Denis Parra and Peter Brusilovsky. 2009. Collaborative filtering for social tagging systems: an experiment with CiteULike. In *Proceedings of the third ACM conference on Recommender systems*. ACM, 237–240.
- [15] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. *Introduction to recommender systems handbook*. Springer.
- [16] Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 253–260.
- [17] Christoph Trattner, Dominik Kowald, and Emanuel Lacic. 2015. TagRec: towards a toolkit for reproducible evaluation and development of tag-based recommender algorithms. *ACM SIGWEB Newsletter* Winter (2015), 3.
- [18] Julita Vassileva. 2008. Toward social learning environments. *IEEE transactions on learning technologies* 1, 4 (2008), 199–214.
- [19] Zhiwen Yu, Xingshe Zhou, and Lei Shu. 2010. Towards a semantic infrastructure for context-aware e-learning. *Multimedia Tools and Applications* 47, 1 (2010).