

Attention Please! A Hybrid Resource Recommender Mimicking Attention-Interpretation Dynamics

Paul Seitlinger
KTI
Graz University of Technology
Graz, Austria
paul.seitlinger@tugraz.at

Ilire Hasani-Mavriqi
KTI
Graz University of Technology
Graz, Austria
ihasani@know-center.at

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

Tobias Ley
CET
Tallinn University
Tallinn, Estonia
tley@tlu.ee

Simone Kopeinik
KTI
Graz University of Technology
Graz, Austria
simone.kopeinik@tugraz.at

Elisabeth Lex
KTI
Graz University of Technology
Graz, Austria
elisabeth.lex@tugraz.at

ABSTRACT

Classic resource recommenders like Collaborative Filtering (CF) treat users as being just another entity, neglecting non-linear user-resource dynamics shaping attention and interpretation. In this paper, we propose a novel hybrid recommendation strategy that refines CF by capturing these dynamics. The evaluation results reveal that our approach substantially improves CF and, depending on the dataset, successfully competes with a computationally much more expensive Matrix Factorization variant.

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

recommender; categorization; connectionism; accuracy; resource recommendations; collaborative filtering; hybrid recommenders; SUSTAIN, attentional focus; decision making; social tagging systems; LDA

1. INTRODUCTION

The Web features a huge amount of data and resources that are potentially relevant and interesting for a user. Users are yet often unable to evaluate all available alternatives due to cognitive limitations of their minds. Recommender systems have been proved as being a valid approach for Web users to cope with the information overload [22] – with Collaborative Filtering (CF) being one of the most successful methods [3]. CF aims to recommend resources to a user

based on the digital traces she leaves behind on the Web, i.e., her interactions with resources and the interactions of other similar users.

Recent advances in the interdisciplinary field of Web Science provide even more comprehensive digital traces of social actions and interactions that can be exploited in recommender research. At least implicitly, recommender research has implemented interesting assumptions about structures and dynamics in Social Information Systems (SIS), such as MovieLens, LastFM or BibSonomy. For instance, by computing matrices or high-dimensional arrays, approaches like CF represent and process SIS as networks or graphs, which relate entities of different quality (e.g., users, resources, time, ratings, tags, etc.) to each other. That way, a compositional view is taken that reminds of a material-semiotic perspective (e.g., [28]), assuming that we gain a deeper understanding of the intention or function of an entity, if we consider the associations it has established with other entities. Put differently, “everything in the social and natural worlds [is regarded] as a continuously generated effect of the webs of relations within which they are located” ([28], p. 142).

However, if we look at the machinery underlying CF, it becomes clear that structurally, the algorithm treats users as being just another entity, such as a tag or a resource. We regard this indifference as a structuralist simplification abstracting from individuals’ complexity. The structuralist stance also runs the risk of neglecting nonlinear, dynamic processes going on between different entities, such as a user’s intentional state (e.g., attentional focus, interpretations, decision making) and resources (articles, movies, etc.) consumed in the past.

It is the main goal of this work to take a closer look at these dynamics and to capture them by means of an appropriate model. Each user develops subjectivity, an idiosyncratic way of perceiving and interpreting things in the world, which manifests itself in particular preferences. Partially, this development evolves through a user’s trajectory in the SIS (e.g., [15]). Every resource that we decide to collect corresponds to a learning episode [11]: Depending on the resource’s features, the episode causes a shift in attention, particularly, in attentional tunings for certain features as well as a shift in mental categories (conceptual clusters), which influences our decision-making (e.g., [29]). The shape that

mental patterns (e.g., attentional tunings and conceptual clusters) acquire, is governed by both, the environment and the current mental state. The acquired pattern in turn orients the user towards particular resources and hence, closes the loop of the environment-user dynamics.

1.1 A Model of Attention-Interpretation Dynamics

The connectionist approach towards human categorization (e.g., [32, 26, 29]) seems to be the method of choice to model these dynamics: Input, hidden and output units are interconnected within a multi-layer network mapping inputs (e.g., resource features) to outputs (e.g., decision to take or leave a resource). Tunings of particular units (e.g., the attentional weight of an input unit) and of the interconnections evolve as the network encounters new exemplars (e.g., resources). A particularly flexible and self-supervising model is SUSTAIN [29], whose number of hidden units (i.e., clusters) is not chosen in advance but incrementally discovered in the course of the learning trajectory. It adapts to the trajectory’s complexity and only recruits a new cluster if the current resource cannot be assimilated with the already existing clusters.

In this work, we train the SUSTAIN network on a user’s history (collected resources in a training set) and then, apply the user-specific network to select new resources from a test set. During training and testing, SUSTAIN encodes the features (LDA topics) of a resource along a set of dimensions at the input layer. The activation of each dimension is gated by an attentional tuning that is learned in the course of the training phase and reflects the importance of the corresponding feature dimension. The hidden layer consists of a set of clusters each representing similar resources encountered in the past. Hence, one cluster corresponds to a user-specific field of interest. SUSTAIN interprets the input, i.e., the pattern of activation along the feature dimensions, by activating each existing cluster according to its similarity to the input. If the activation of the most similar cluster is below a particular threshold, SUSTAIN forms a new cluster to simulate the mental process of accommodation. In the test phase, where the set of recruited clusters does no longer change, the classification decision (in our case, the decision of choosing a given resource or not) is a function of the activation of the most similar cluster.

SUSTAIN can be applied to do more justice to a user’s way of interacting within an SIS, e.g., in terms of the features (topics) the user attends to and of the semantic interpretation of the resources. For instance, consider a user exhibiting a broad attentional focus because the collected resources cover a broad range of topics. Modeling that user’s trajectory by means of SUSTAIN will yield a network with a uniform distribution of attentional tunings (i.e., for different topics) and a number of distinct clusters since the resources fall inside different regions of the topic space. On the other hand, consider a user with a narrow attentional focus that corresponds to a list of resources addressing similar topics. The result will be a peaky attentional distribution and a small number of clusters.

Since SUSTAIN is able to differentiate between users with respect to attention and interpretation, we assume that it can be applied to anticipate user-specific preferences and decisions on resources. In this work, we therefore introduce a resource recommender that draws on SUSTAIN to model

a user’s traces (e.g., posts) in form of a connectionist network representing the user’s attentional focus and semantic clusters. Given the traces provide sufficient information for training the network, we assume that a recommender equipped with SUSTAIN can be applied to simulate a user’s decision making on a set of resources and therefore, improve the recommender’s accuracy. For a first proof of concept, we proceed in two steps: First, we make use of CF to gather a more restricted set of resources. Second, SUSTAIN simulates how the user interacts with each of the resources and simulates the resulting decision of the user to take or leave the resource. These two steps constitute our novel hybrid resource recommender approach termed SUSTAIN+CF_U (see Section 3). In brief, our research question of this work is:

RQ: *Do resource recommendations become more accurate if a set of resources identified by CF is processed by SUSTAIN to simulate user-specific attentional and conceptual processes?*

The remainder of this paper is organized as follows: In Section 2, we consider related work that has inspired our own approach, whose algorithm is described in Section 3. In Section 4, we describe our methodology applied to compare the performance of our SUSTAIN+CF_U approach with several baseline algorithms. In Sections 5 and 6, we then report our evaluation results and provide a conclusion and outlook, respectively.

2. RELATED WORK

At the moment, we identify three main research disciplines that are related to our work.

Collaborative Filtering Extensions: One of our previous studies in this field [27], introduces the so-called *Collaborative Item Ranking Using Tag and Time Information (CIRTT)* approach, which extends CF in social tagging systems by incorporating tag and time information. This approach combines user-based and item-based CF with the information of tag frequency and recency by applying the base-level learning (BLL) equation coming from human memory theory.

An extensive survey on CF was recently conducted by [37]. In this survey, the authors classify CF approaches based on the type of information that is processed and the type of paradigm applied. Furthermore, CF extensions are defined as approaches which enrich classic CF algorithms with valued side information on users and resources. Analogous categorization of CF studies is performed in [2] as well. Additionally, these studies have identified challenges that are crucial to the future research on CF. In this context, authors state the fact that there is a lack of studies which address issues on recommender systems from the psychological perspective. To the best of our knowledge, there has been no remarkable endeavors which tend to combine the implementation of a dynamic and connectionist model of human cognition, such as SUSTAIN, with existing CF algorithms.

Recommender Systems and User Modeling: The work of [10] distinguishes between recommender systems that provide non-personalized and personalized recommendations. Whereas, non-personalized recommender systems are not based on user models, personalized ones choose resources by considering the user profile (e.g. previous user interactions or user preferences). Considerable amount of techniques have been proposed to design the user model in terms of resource recommendations [21] [8]. Among them,

some approaches aim to provide dynamically adapted personalized recommendations to users [13].

A specific research topic, which is increasingly gaining popularity, is human decision making in recommender systems [7]. The work presented in [9] systematically analyzes recommender systems as decision support systems based on the nature of users’ goals and the dynamic characteristics of the resource space (e.g., availability of resources). However, there is still lack of research on investigating user decision processes at a detailed level and considering integrating scientific facts from psychology. Thus, we consider that our proposed approach contributes to this line of research.

Long Tail Recommendations and User Serendipity: Another concept within the area of recommender systems, that needs to be more extensively considered in the future, is termed as long tail recommendations. Basically, the long tail refers to the resources that have low popularity [37]. It is of a huge interest to show how recommendations of these long tail resources can impact user satisfaction. Furthermore, it is important to investigate if additional revenue can be generated by the recommender systems from the long tail resources [37] [39] [36].

In this line of research, various solutions have been proposed to overcome the problem of over-specialization and concentration-bias in recommender systems [1]. The problem of concentration-bias becomes evident since traditional CF algorithms recommend resources based on the users’ previous history of activities. Hence, resources with the most occurrences in this history are typically repeatedly recommended to users causing narrowing the choices, by excluding other resources which might be of interest. Additionally, recommending resources based on user’s previous activities or preferences yields to over-specialization of recommendations. However, balancing between information overload and facilitating users to explore new horizons by recommending serendipitous choices is not tackled within the scope of this work. It still remains a challenge to be considered in our future research.

3. APPROACH

In this section, we introduce a novel hybrid resource recommender approach termed SUSTAIN+CF_U, which further personalizes and improves user-based Collaborative Filtering (CF_U). SUSTAIN [29], briefly described in the introductory section, is a computational model of human category learning. It builds upon the assumption that things can be described by patterns of correlated features, that form n -dimensional clusters, and attentional tunings representing each dimension’s relative importance. To apply this theory when recommending Web resources, we use 500 LDA topics derived from the tag assignments as the n features [17] (see Section 4.1). On the basis of the resources a user has bookmarked in the past (i.e., the training set of a user), each user’s personal attentional tunings and cluster representations are created in a training phase. Following an unsupervised learning procedure, we start simple, with one cluster and expand the number of clusters if necessary. Please note, that all SUSTAIN-specific parameter settings are adopted from [29] (see Table 1).

For each resource in the training set of user u , we start by calculating the distance μ_{ij} to cluster j at dimension i as described in equation (1):

| Function | Symbol | Value |
|---------------------|---------|-------|
| Attentional focus | r | 9.998 |
| Cluster competition | β | 6.396 |
| Learning rate | η | .096 |
| Threshold | η | .5 |

Table 1: SUSTAIN’s best fitting parameters for unsupervised learning as suggested in [29].

$$\mu_{ij} = I^{pos_i} - H_j^{pos_i} \quad (1)$$

where I is the n -dimensional input vector, which represents the topics of this resource, and vector H_j is cluster j ’s position in the n -dimensional feature space, which holds a value for each topic and is initially set to $\vec{0}$. In the next step, the distance μ_{ij} is used to calculate the activation value H_j^{act} of the j^{th} cluster by equation (2):

$$H_j^{act} = \frac{\sum_{i=1}^n (\lambda_i)^r e^{-\lambda_i \mu_{ij}}}{\sum_{i=1}^n (\lambda_i)^r} \quad (2)$$

where λ_i represents the tuning (weight) of dimension i and acts as a multiplier on i in calculating the activation. Initially, vector λ is set to $\vec{1}$ and evolves during the training phase according to equation (4) calculated at the end of every training iteration (i.e., after including a resource). r , which is set to 9.998, is an attentional parameter that accentuates the effect of λ_i : if $r = 0$, all dimensions are weighted equally. Clusters compete with one another, thus weakening the most activated one (i.e., $\max(H_j^{act})$). This form of lateral inhibition is calculated by equation (3), which gives an activation value H_m^{out} for the winning cluster m :

$$H_m^{out} = \frac{(H_m^{act})^\beta}{\sum_{i=1}^n (H_i^{act})^\beta} H_m^{act} \quad (3)$$

where β is the lateral inhibition parameter and is set to 6.396. If the activation value H_m^{out} of the most activated (i.e., winning) cluster is beneath a given threshold $\tau = .5$, a new cluster is created, representing the topics of the currently processed resource. At the end of an iteration, the tunings of vector λ are updated given by equation (4):

$$\Delta \lambda_i = \eta^{-\lambda_i \mu_{im}} (1 - \lambda_i \mu_{im}) \quad (4)$$

where j indexes the winning cluster and the learning rate η is set to .096. In a final step, the position vector of the winning cluster, which holds a value for each of the n topics, is recalculated as described by equation (5):

$$\Delta H_m^{pos_i} = \eta (I^{pos_i} - H_m^{pos_i}) \quad (5)$$

The training phase is completed when steps (1) to (5) were subsequently processed for every resource in a user’s training set. For each user, this results in a particular vector of attentional tunings λ and a set of j cluster vectors H_j .

We make use of this user-specific pattern to calculate personalized resource recommendations by incorporating CF_U into our hybrid approach SUSTAIN+CF_U. Thus, first, we determine the top-100 resources identified by CF_U (see Section 4.3) as a candidate set C_u of potentially relevant resources for the target user u . Then, for each candidate c in C_u , we calculate H_m^{out} following equations (1) to (3).

$H_m^{out}(c)$ can then be added to $CF_U(c)$ as shown in equation (6) in order to determine the set of k recommended resources $RecRes(u)$ for user u :

$$RecRes(u) = \arg \max_{c \in C_u} \underbrace{(\alpha H_m^{out}(c) + (1 - \alpha) CF_U(c))}_{SUSTAIN + CF_U} \quad (6)$$

where α can be used to inversely weigh the two components of our hybrid approach. For now, we set α to .5 in order to equally weight SUSTAIN and CF_U .

Our approach as well as the baseline algorithms described in Section 4.3 (but for WRFM) and the evaluation method described in Section 4.2 are completely implemented in Java within our *TagRec* recommender benchmarking framework [24], which is freely available via GitHub¹.

4. METHODOLOGY

In this section we describe the methodology used in this work, including the datasets, the evaluation method and metrics and the baseline algorithms.

4.1 Datasets

For the evaluation of our approach, we used datasets gathered from well-known social tagging systems. We focused on social tagging systems for our study not only because their datasets are freely-available for scientific purposes but also because tagging data can be easily utilized to derive semantic topics for resources from it [17] by means of Latent Dirichlet Allocation (LDA) [25]. This initial step was necessary since our datasets did not contain categories or topics for the resources (e.g., like Wikipedia categories, see [35]) that are the basis for our approach. Thus, we chose LDA to simulate this external categorization.

LDA is a probability model that helps to find latent semantic topics for documents (i.e., resources). In the case of tagging data, the model takes the assigned tags of the resources as input and returns the identified topic distributions for each resource. We implemented LDA using the Java framework Mallet² with Gibbs sampling and 2000 iterations as suggested in the framework documentation and by related work (e.g., [25]). Moreover, we set the number of latent topics to 500 (see also [23, 35]) and considered only identified topics for resources that have a minimum probability value of .01 in order to reduce noise and to meaningfully limit the number of assigned topics.

We focused on the social bookmark and publication sharing system BibSonomy³ (2013-07-01), the citation sharing system CiteULike⁴ (2013-03-10) and the social bookmarking system Delicious⁵ (2011-05-01) to test our approach in three different settings that differ in their dataset sizes. To reduce computational effort, we randomly selected 20% of the CiteULike user profiles [16] (the other datasets were processed in full size). We did not use a p -core pruning approach to avoid a biased evaluation (see [12]) but excluded all posts assigned to unique resources, i.e., resources that have only been bookmarked once (see [31]). The statistics of the full

| Dataset | Type | $ P $ | $ U $ | $ R $ | $ T $ |
|-----------|--------|---------|--------|---------|---------|
| BibSonomy | Full | 400,983 | 5,488 | 346,444 | 103,503 |
| | Sample | 82,539 | 2,437 | 28,000 | 30,919 |
| CiteULike | Full | 753,139 | 16,645 | 690,126 | 238,109 |
| | Sample | 105,333 | 7,182 | 42,320 | 46,060 |
| Delicious | Full | 104,799 | 1,867 | 69,223 | 40,897 |
| | Sample | 59,651 | 1,819 | 24,075 | 23,984 |

Table 2: Properties of the full datasets as well as the used dataset samples, where $|P|$ is the number of posts, $|U|$ is the number of users, $|R|$ is the number of resources and $|T|$ is the number of tags.

datasets as well as our used dataset samples after the exclusion of the posts assigned to unique resources are shown in Table 2.

4.2 Evaluation Method and Metrics

In order to evaluate our algorithm and to follow common practice in recommender systems research (e.g., [20, 40]), we split our datasets into training and test sets. Therefore, we followed the method described in [27] to retain the chronological order of the posts. This also simulates well a real-world environment, where future interactions are tried to be predicted based on interactions in the past [5]. Hence, for each user we used her 20% most recent posts for testing and the rest for training.

To finally determine the performance of our approach as well as of the baseline methods, we compared the top-20 recommended resources determined by each algorithm for a user with the relevant resources in the test set using a variety of well-known evaluation metrics [30, 18] in recommender systems research. In particular, we report Normalized Discounted Cumulative Gain (nDCG@20), Mean Average Precision (MAP@20), Recall (R@20) and Precision (P@20). Moreover we show the performance of the algorithms for different numbers of recommended resources ($k = 1 - 20$) by means of Precision / Recall plots.

4.3 Baseline Algorithms

We selected a set of well-known resource recommender baseline algorithms in order to determine the performance of our novel approach in relation to these approaches. Hence, we have not only chosen algorithms that are similar to our approach in terms of their processing steps (CF_U and CB_T) but also current state-of-the-art methods for personalized resource recommendations (CF_R and WRFM) along with a simple unpersonalized approach (MP).

Most Popular (MP): The most simple method we compare our algorithm to, is the *Most Popular (MP)* approach that ranks the resources by their total frequency in all posts [30]. In contrast to the other chosen baselines, the MP approach is non-personalized and thus, recommends for any user the same set of resources.

User-based Collaborative Filtering (CF_U): *User-based Collaborative Filtering (CF_U)* typically consists of two steps. First, the most similar users (the so-called k nearest neighbors) for a target user are found based on a specific similarity measure and second, the resources of these neighbors, that are new to the target user, are suggested. This procedure is based on the idea, that if two users had a similar taste in the past, they will probably share the same

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

²<http://mallet.cs.umass.edu/>

³<http://www.kde.cs.uni-kassel.de/bibsonomy/dumps/>

⁴<http://www.citeulike.org/faq/data.adp>

⁵<http://files.grouplens.org/datasets/hetrec2011/hetrec2011-delicious-2k.zip>

| Dataset | Metric | MP | CF _R | CB _T | WRMF | CF _U | SUSTAIN+CF _U |
|-----------|---------|-------|-----------------|-----------------|--------------|-----------------|-------------------------|
| BibSonomy | nDCG@20 | .0142 | .0569 | .0401 | .0491 | .0594 | .0665 |
| | MAP@20 | .0057 | .0425 | .0211 | .0357 | .0429 | .0492 |
| | R@20 | .0204 | .0803 | .0679 | .0751 | .0780 | .0859 |
| | P@20 | .0099 | .0223 | .0272 | .0132 | .0269 | .0290 |
| CiteULike | nDCG@20 | .0064 | .1006 | .0376 | .0411 | .0753 | .0857 |
| | MAP@20 | .0031 | .0699 | .0170 | .0210 | .0468 | .0555 |
| | R@20 | .0090 | .1332 | .0697 | .0658 | .1149 | .1251 |
| | P@20 | .0023 | .0289 | .0174 | .0218 | .0257 | .0280 |
| Delicious | nDCG@20 | .0038 | .1148 | .0335 | .1951 | .1327 | .1633 |
| | MAP@20 | .0011 | .0907 | .0134 | .1576 | .0949 | .1156 |
| | R@20 | .0071 | .1333 | .0447 | .2216 | .1662 | .1946 |
| | P@20 | .0017 | .0512 | .0173 | .1229 | .0843 | .0982 |

Table 3: nDCG@20, MAP@20, R@20 and P@20 estimates for BibSonomy, CiteULike and Delicious showing that our proposed approach SUSTAIN+CF_U not only outperforms CF_U in all settings but also is able to successfully compete with WRMF (Note: highest accuracy values per dataset over all algorithms are highlighted in bold).

taste in the future and thus, will like the same resources [34]. We calculated the user-similarities based on the binary user-resource matrix and using the cosine-similarity measure (see [40]). Moreover, we set the neighborhood size k to 20 as suggested in [16] for CF_U in social tagging systems.

Resource-based Collaborative Filtering (CF_R): In contrast to CF_U, *Resource-based Collaborative Filtering (CF_R)* (also known as Item-based CF), identifies potentially interesting resources for a user via computing similarities between resources instead of similarities between users. Hence, this approach processes the resources a user has bookmarked in the past in order to find similar resources to be recommended [33]. As in the case of CF_U, we calculated similarities based on the binary user-resource matrix using cosine similarity and focused on a resource-neighborhood size k of 20 [40, 16].

Content-based Filtering using Topics (CB_T): Content-based filtering (CB) methods recommend resources to users based on a comparison between the resource content and the user profile [4]. Hence, this approach does not need to calculate similarities between users or resources (as done in CF methods) but directly tries to map resources and users. We implemented this method in the form of *Content-based Filtering using Topics (CB_T)* since topics are the only content-based features available in our social tagging datasets (see Section 4.1). The similarity between the topic vector of a user and a resource has been calculated using the cosine similarity measure.

Weighted Regularized Matrix Factorization (WRMF): *WRMF* is a model-based recommender method for implicit data (e.g., posts) based on the state-of-the-art Matrix Factorization (MF) technique. MF factorizes the binary user-resource matrix into latent user- and resource-factors, that represent these entities, in a common space. This representation is used to map resources and users and thus, find resources to be recommended for a specific user. WRMF defines this task as a regularized least-squares problem based on a weighting matrix, which differentiates between observed and unobserved activities in the data [19]. The results for WRFM presented in Section 5 have been calculated using the MyMediaLite 3.10 framework⁶ (2013-

09-23) with 500 latent factors, 100 iterations and a regularization value of .001.

5. RESULTS

A glance at the values in Table 3 and the Precision / Recall plots in Figure 1 reveals that the simplest baseline algorithm, the unpersonalized MP approach, results in very low estimates of accuracy. This indicates the difficulty of the task of predicting users’ resource choices in our three datasets gathered from the social tagging systems BibSonomy, CiteULike and Delicious. However, across all datasets, the remaining personalized algorithms reach larger estimates and therefore, are successful in explaining a substantial amount of variance in user behavior.

Referring to our research question RQ stated in Section 1, as to whether SUSTAIN+CF_U can improve CF_U in terms of recommender accuracy, our hybrid approach outperforms CF_U with respect to all metrics applied and across all three datasets. For instance, in the Precision / Recall plots in Figure 1, we can see no overlap between the two corresponding curves – with SUSTAIN+CF_U always reaching higher values than CF_U. Moreover, especially the results of the ranking-dependent metric nDCG@20 in Table 3 shows a large difference between SUSTAIN+CF_U and CF_U, indicating that our approach can be used to successfully re-rank the candidate resources identified by CF_U in a more personalized manner. We attribute this to the fact that there is no specific ranking of the resources of a neighbor in user-based CF applied to implicit data. So it can happen that the whole recommended list contains of resources of the nearest neighbor without any ranking, which is tried to be avoided in our hybrid approach. Summed up, we can answer our RQ in the affirmative.

Interestingly, the results also clearly show that the performance of the algorithms strongly varies across BibSonomy, CiteULike and Delicious and that in each of the three datasets a different algorithm performs best. For instance, in the case of CiteULike, the best results are reached by CF_R. We relate this to the fact that the average topic similarity per user in CiteULike (18.9%) is much higher than in BibSonomy (7.7%) and Delicious (4.5%), indicating a more thematically consistent resource search behavior. The higher consistency is in advantage for predictions that are based on resources collected in the past, such as CF_R-

⁶<http://www.mymedialite.net/>

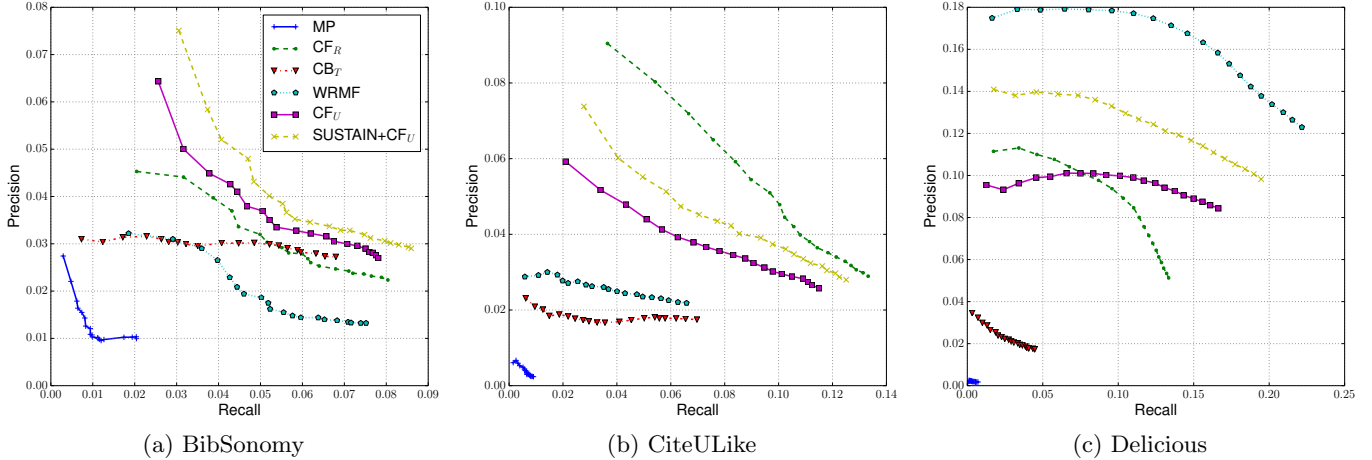


Figure 1: Precision / Recall plots for BibSonomy, CiteULike and Delicious showing the recommender accuracy of our approach SUSTAIN+CF_U in comparison to the baseline methods for $k = 1 - 20$ recommended resources. The results indicate that our approach provides higher Precision and Recall estimates than CF_U for each k and in all three datasets. In the case of BibSonomy SUSTAIN+CF_U even outperforms all baseline methods, including CF_R and WRMF.

based predictions. In case of Delicious, the users in the dataset are chosen through a so-called mutual-fan crawling strategy (see [6]) and thus, are not independent from each other. This is conducive to methods that capture relations between users with common resources by means of high-dimensional arrays, such as WRMF. However, compared to the other algorithms, especially to CF_R and WRMF, SUSTAIN+CF_U results in relatively robust estimates since our provides fairly good results in all three datasets. Finally, SUSTAIN+CF_U even outperforms all other algorithms in case of BibSonomy.

6. CONCLUSIONS AND FUTURE WORK

With this work we show that a connectionist model of human category learning applied to mimic a user’s attentional focus and interpretation helps to improve user-based CF predictions. We attribute this improvement to the well established, cognitive plausibility of SUSTAIN [29] by which in our case a larger amount of variance in a user’s decision making on a given set of resources can be explained: Reconstructing the user history in form of an iteratively trained network with history-specific patterns of attentional tunings and clusters, does more justice to a user’s individuality than a CF-based representation of user-resource relations. Additionally, we observe that our hybrid SUSTAIN+CF_U model is more robust in terms of accuracy estimates but also less complex in terms of computational effort than WRMF.

In future, we aim at improving and further validating our model. First, we are working on a variant that is independent of CF_U and searches for user-specific recommendations only by means of the correspondingly trained SUSTAIN network. Second, we will make use of the network’s sensitivity for a user’s mental state to realize a more dynamic recommendation logic. In particular, based on creative cognition research (e.g., [14]), we assume a broader attentional focus to be associated with a stronger orientation toward novel or more diverse resources. Recommendations should become more accurate and diverse if this kind of association turns

out to be reliable and is integrated into the algorithm. Furthermore, we will prove the claim that our approach is less complex in terms of computational costs than WRMF by means of a detailed runtime and memory analysis.

With respect to evaluation, we regard algorithmic models, such as SUSTAIN, not only as tools to improve recommendations but also to derive specific research questions in the field of Web Science. For instance, we can raise the question whether the extent of a user’s curiosity is related to her ability to discover high quality resources before other users. One way to tackle this question would be to draw on the established SPEAR algorithm [38], differentiating between discoverers and followers and relating its scores with SUSTAIN-specific scores of curiosity.

We conclude that our attempt to keep the translation from theory into technology as direct as possible holds advantages for both technical and conceptual aspects of recommender research. By applying computational models of human cognition, we can improve the performance of existing recommender mechanisms and at the same, gain a deeper understanding of fine-grained level dynamics in a Social Information System.

7. ACKNOWLEDGMENTS

The authors would like to thank Emanuel Lacic and Sebastian Dennerlein for many valuable comments on this work. This work is supported by the Know-Center, the EU funded projects Learning Layers (Grant Agreement 318209) and weSPOT (Grant Agreement 318499) and the Austrian Science Fund (FWF): P 25593-G22. 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).

8. REFERENCES

- [1] P. Adamopoulos and A. Tuzhilin. On over-specialization and concentration bias of recommendations. In *Proc. of RecSys'14*, RecSys '14, pages 153–160, New York, NY, USA, 2014. ACM.
- [2] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, 2005.
- [3] A. Bar, L. Rokach, G. Shani, B. Shapira, and A. Schlar. Improving simple collaborative filtering models using ensemble methods. In *Multiple Classifier Systems*, pages 1–12. Springer, 2013.
- [4] J. Basilico and T. Hofmann. Unifying collaborative and content-based filtering. In *Proc. of ICML'04*, page 9. ACM, 2004.
- [5] P. G. Campos, F. Díez, and I. Cantador. Time-aware recommender systems: a comprehensive survey and analysis of existing evaluation protocols. *User Modeling and User-Adapted Interaction*, pages 1–53, 2013.
- [6] I. Cantador, P. Brusilovsky, and T. Kuflik. 2nd workshop on information heterogeneity and fusion in recommender systems (hetrec 2011). In *Proc. of RecSys'11*, RecSys 2011, New York, NY, USA, 2011. ACM.
- [7] L. Chen, M. de Gemmis, A. Felfernig, P. Lops, F. Ricci, and G. Semeraro. Human decision making and recommender systems. *ACM Trans. Interact. Intell. Syst.*, 3(3):17:1–17:7, Oct. 2013.
- [8] B. Coleho, C. Martins, and A. Almeida. Web intelligence in tourism: User modeling and recommender system. In *Proc. of WI-IAT '10*, pages 619–622, Washington, DC, USA, 2010. IEEE Computer Society.
- [9] P. Cremonesi, A. Donatucci, F. Garzotto, and R. Turrin. Decision-making in recommender systems: The role of user's goals and bounded resources. In *Proc. of Decisions@RecSys'12 Workshop*, volume 893 of *CEUR Workshop Proceedings*, pages 1–7. CEUR-WS.org, 2012.
- [10] P. Cremonesi, F. Garzotto, and R. Turrin. Investigating the persuasion potential of recommender systems from a quality perspective: An empirical study. *ACM Trans. Interact. Intell. Syst.*, 2(2):11:1–11:41, June 2012.
- [11] S. Dennerlein, M. Rella, V. Tomberg, D. Theiler, T. Treasure-Jones, M. Kerr, T. Ley, M. Al-Smadi, and C. Trattner. Making sense of bits and pieces: A sensemaking tool for informal workplace learning. In *Open Learning and Teaching in Educational Communities*, pages 391–397. Springer, 2014.
- [12] S. Doerfel and R. Jäschke. An analysis of tag-recommender evaluation procedures. In *Proc. of RecSys'13*, pages 343–346, New York, NY, USA, 2013. ACM.
- [13] S. Dooms. Dynamic generation of personalized hybrid recommender systems. In *Proc. of RecSys'13*, RecSys '13, pages 443–446, New York, NY, USA, 2013. ACM.
- [14] R. A. Finke, T. B. Ward, and S. M. Smith. Creative cognition: Theory, research, and applications. 1992.
- [15] W.-T. Fu and W. Dong. Collaborative indexing and knowledge exploration: A social learning model. *IEEE Intelligent Systems*, 27(1):39–46, 2012.
- [16] J. Gemmell, T. Schimoler, M. Ramezani, L. Christiansen, and B. Mobasher. Improving folkrank with item-based collaborative filtering. *Recommender Systems & the Social Web*, 2009.
- [17] T. L. Griffiths, M. Steyvers, J. B. Tenenbaum, et al. Topics in semantic representation. *Psychological review*, 114(2):211, 2007.
- [18] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22(1):5–53, Jan. 2004.
- [19] Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*, pages 263–272. IEEE, 2008.
- [20] C.-L. Huang, P.-H. Yeh, C.-W. Lin, and D.-C. Wu. Utilizing user tag-based interests in recommender systems for social resource sharing websites. *Knowledge-Based Systems*, 56:86–96, 2014.
- [21] G. Jawaheer, P. Weller, and P. Kostkova. Modeling user preferences in recommender systems: A classification framework for explicit and implicit user feedback. *ACM Trans. Interact. Intell. Syst.*, 4(2):8:1–8:26, June 2014.
- [22] P. B. Kantor, L. Rokach, F. Ricci, and B. Shapira. *Recommender systems handbook*. Springer, 2011.
- [23] W. Kintsch and P. Mangalath. The construction of meaning. *Topics in Cognitive Science*, 3(2):346–370, 2011.
- [24] D. Kowald, E. Lacic, and C. Trattner. Tagrec: Towards a standardized tag recommender benchmarking framework. In *Proc. of HT'14*, New York, NY, USA, 2014. ACM.
- [25] R. Krestel, P. Fankhauser, and W. Nejdl. Latent dirichlet allocation for tag recommendation. In *Proc. of Recsys'09*, pages 61–68. ACM, 2009.
- [26] J. K. Kruschke et al. Alcove: An exemplar-based connectionist model of category learning. *Psychological review*, 99(1):22–44, 1992.
- [27] E. Lacic, D. Kowald, P. Seitlinger, C. Trattner, and D. Parra. Recommending items in social tagging systems using tag and time information. In *Proc. of the Social Personalization Workshop colocated with HT'14*, 2014.
- [28] J. Law. Actor network theory and material semiotics. *The new Blackwell companion to social theory*, pages 141–158, 2009.
- [29] B. C. Love, D. L. Medin, and T. M. Gureckis. Sustain: a network model of category learning. *Psychological review*, 111(2):309, 2004.
- [30] D. Parra and S. Sahebi. Recommender systems : Sources of knowledge and evaluation metrics. In *Advanced Techniques in Web Intelligence-2: Web User Browsing Behaviour and Preference Analysis*, pages 149–175. Springer-Verlag, 2013.
- [31] D. Parra-Santander and P. Brusilovsky. Improving collaborative filtering in social tagging systems for the recommendation of scientific articles. In *Proc. of WIAT'10*, volume 1, pages 136–142. IEEE, 2010.
- [32] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Cognitive modeling*, 1988.
- [33] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *Proc. of WWW'01*, pages 285–295. ACM, 2001.
- [34] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen. Collaborative filtering recommender systems. In *The adaptive web*, pages 291–324. Springer, 2007.
- [35] P. Seitlinger, D. Kowald, C. Trattner, and T. Ley. Recommending tags with a model of human categorization. In *Proc. of CIKM'13*, pages 2381–2386, New York, NY, USA, 2013. ACM.
- [36] L. Shi. Trading-off among accuracy, similarity, diversity, and long-tail: A graph-based recommendation approach. In *Proc. of RecSys'13*, pages 57–64, New York, NY, USA, 2013. ACM.
- [37] Y. Shi, M. Larson, and A. Hanjalic. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Comput. Surv.*, 47(1):3:1–3:45, May 2014.
- [38] C.-m. A. Yeung, M. G. Noll, N. Gibbins, C. Meinel, and N. Shadbolt. Spear: Spamming-resistant expertise analysis and ranking in collaborative tagging systems. *Computational Intelligence*, 27(3):458–488, 2011.
- [39] H. Yin, B. Cui, J. Li, J. Yao, and C. Chen. Challenging the long tail recommendation. *Proc. VLDB Endow.*, 5(9):896–907, May 2012.
- [40] N. Zheng and Q. Li. A recommender system based on tag and time information for social tagging systems. *Expert Systems with Applications*, 38(4):4575–4587, 2011.