Recommendation in the Social Web

Robin Burke and Jonathan Gemmel
Web Intelligence Laboratory
DePaul University
Chicago, IL, USA

Andreas Hotho and Robert Jäschke
Knowledge and Data Engineering
University of Kassel
Kassel, Germany

Abstract

Recommender systems are a means of personalizing the presentation of information to ensure that users see the items most relevant to them. The social web has added new dimensions to the way people interact on the Internet, placing the emphasis more heavily on user-generated content. Users in social networks create photos, videos and other artifacts, collaborate with other users, socialize with their friends and share their opinions online. This outpouring of material has brought increased attention to recommender systems, as a means of managing this vast universe of content. At the same time, the diversity and complexity of the data has meant new challenges for researchers in recommendation. This article describes the nature of recommendation research in social web applications and provides some illustrative examples of current research directions and techniques.

It is difficult to overstate the impact of the social web. This new breed of social applications is reshaping nearly every human activity from the way people watch movies to how they overthrow governments. Facebook allows its members to maintain friendships whether they live next door or on another continent. With Twitter, users from celebrities to ordinary folks can launch their 140 character messages out to a diverse horde of "followers." Flickr and YouTube users upload their personal media to share with the world, while Wikipedia editors collaborate on the world's largest encyclopedia.

These applications vary in the manner of the user interaction. Yet they all share some fundamental characteristics. First, the focus is often partly (if not entirely) on the users themselves, and especially on their assocations and interactions with others. Second, many of these applications facilitate the organization and creation of online content. Finally, social applications often allow users to respond to posted content in a variety of ways, embedding each item in an evolving social network.

The complex information space generated by the social web offers a rich and dynamic environment for users to share information, discover new content and meet new people. However, the success of the social web has made some of these benefits difficult to realize, due to the vast amount of information available. Twitter has 50 million tweets per day; which should you read? Facebook has 600 million active users; who should be your friend? Every minute more than 24 hours of video is uploaded to YouTube; what should you watch? Recommender systems have the potential to filter these oceans of data making up the social web, and provide a personalized view that matches each users' interest.

The types of recommender systems most commonly studied by researchers and discussed elsewhere in this special issue are not always a great fit for the task of recommendation in the social web. Data used to recommend products in an e-commerce setting, for example, represent a user-item relationship: a user might give a movie three stars, or a database might show which products a user has purchased. We can think of this as two-dimensional data, and much of recommender systems research has been focused on algorithms appropriate to data in this form.

The data associated with the social web is radically multi-dimensional in comparison. There are many more kinds of recommendation that can be performed. The applications that focus on building social networks, like Facebook and LinkedIn, need engines that recommend users rather than items. Applications that permit the annotation of online items like Flickr and LastFM must recommend items and tags to annotate them. The types of recommendation are so numerous and the underlying data so varied from application to application that no single recommendation strategy could hope to offer a definitive solution.

This article describes the landscape of the social web, sampling some of the components common across its many applications, and discusses some of the recommendation tasks that naturally arise from this landscape. We then sample a few of the current techniques being employed to address these challenges.

Landscape of the Social Web

The social web can be considered as the aggregation of online interactions among users. These interactions can take a myriad of forms, some of which are shown in Figure 1. Some interactions are between individuals while others revolve around online content or how that content is valued. In this section, we examine some of the core entities found in the social web and their characteristics with respect to recommendation.

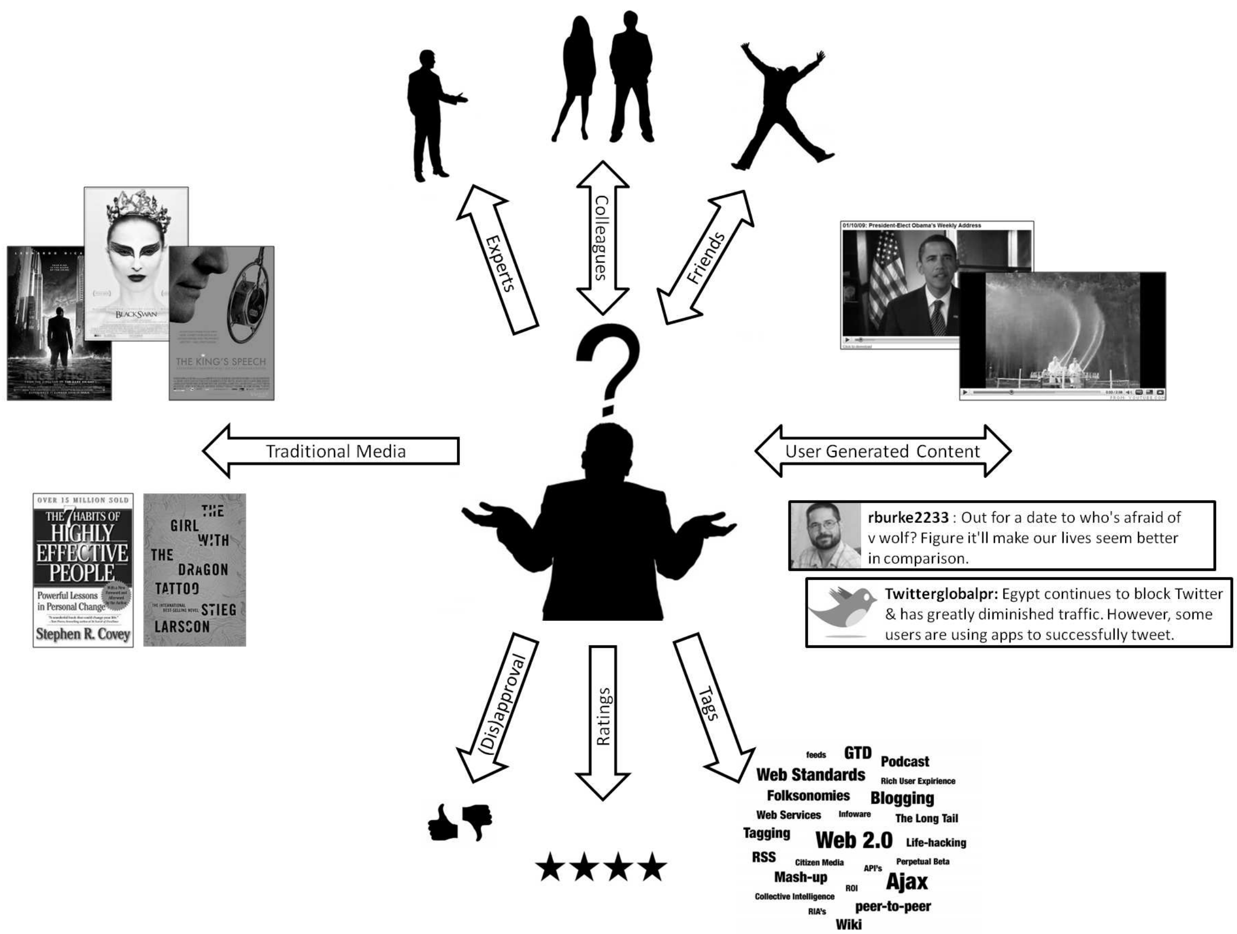

Figure 1: Landscape of the Social Web: A user has many options on how to with connect other users, what types of content to consume, and how to respond to this content.

People

Research on the World Wide Web often models that system as an information graph, with web pages linking to each other in complex patterns. In this model, the authors do not participate and indeed, anonymity and difficulty in establishing authorship are two hallmarks of the web as we know it. The social web, on the other hand, lends itself to graphs in which users are connected to each other, either directly through social links of various types or indirectly through connections to content.

The interactions between a user and his online social contacts can vary as dramatically as they do in real world environments. Interactions can be directional, when they are formed through subscription arrangements, as in Twitter, where the user being followed has no control over who chooses to be follower. They can also be bi-directional – as in the now famous "friend" relationship on Facebook. Interactions may involve groups of different sizes. Sometimes a user is posting content to a closed group of friends; sometimes, to the whole world. These actions all require different interpretations.

A basic recommendation operation in the social web is the recommendation of other users. Friend recommendation, for example, is a common operation that can directly leverage the social graph, as in Facebook's built-in friend finder that locates individuals with overlapping social circles. Still, because of the nature of the social environment, user recommendation has multiple interpretations. A user seeking a recommendation might be looking for a friend, but she might also be seeking a colleague, a potential client, a romantic prospect, a like-minded hobbyist, or any other type of social contact. She might be seeking someone she already knows or someone new. In some applications, especially corporate intranets, the recommendation of indviduals with particular skills or expertise may be an important task. Such an expert is unlikely to be found in the user's immediate social circle.

Items

The social web presents a vast array of items that can be recommended. One reason that people come to social web sites is to hear about new items of the type that recommender systems have typically presented: movies, music, restaurants. However, in the social web, there will often be multiple routes into a given item. A music track might be accessed through the artist, through a tag, through an album, through a friend's "favorites" list, etc. The multitude of links and interactions surrounding an item provides additional clues about its meaning and significance, but also poses a challenge for the recommender.

In addition to exogenous media items, a social web site will also contain user-generated content, which is also fair game for recommendation. User-generated items, such as blog entries, status updates, and other data, have the benefit of being closely embedded in the networks of interaction that a social web site hosts. Often, however, they are not intended to stand alone and can sometimes only be interpreted within that context. For example, the Twitter post about "Who's Afraid of Virginia Woolfe?" in Figure 1 makes much more sense if you know that the poster is a college professor, like the male lead in the play.

Recommendations of user-generated content must also take context into account. Consider a user who just viewed photographs annotated with "wedding" and "New York". Would such a user be interested in seeing all photos with these tags? Most likely, the user interest is related to the particular individuals getting married, and pictures of strangers' New York weddings would not be welcome recommendations.

Valuations

Users in social web sites have a variety of options for reacting to and expressing their opinions about contents of the site. They can tag, post, tweet, rate, or write reviews. These reactions can be read, responded to, or forwarded to others. Substantial valuations like reviews may become items of recommendation in their own right. Many users prefer a peer review to a professionally-authored one, especially if the author has some commonality with themselves.

One of the most well-studied recommendation applications in the social web is the social annotation or tagging system. In these social web systems, the main valuation provided by users is in the form of tags, short text labels in an unconstrained vocabulary. Some examples include the music sharing site Last.fm and the citation sharing site BibSonomy. Part of the appeal of these systems is that they show potential for distributing the knowledge engineering involved in creating representations of items: the idea of a "folksonomy" as an alternative to an officially-created "taxonomy.".

In social annotation systems, a common application of recommendation is to assist users in labeling items through the recommendation of tags. If properly implemented, tag recommendation reduces the cognitive load of tagging, encouraging users to do more of it, and also reduces the amount of noise in the tag vocabulary by reducing redundancies like "NewYork", "New-York", "New_York", etc.[1]

Recommendation Tasks

The complexity of recommendation in the social web is not solely a function of the variety of possible items to be recommended. Its multi-dimensional nature means that a wide variety of recommendation tasks and modalities can be supported.

We can distinguish recommendation tasks by their input and output, and by the semantics of the recommendation operation. The wide variety of outputs has already been noted. The essential input to any recommender system is the profile of the user for whom the recommendation is sought. However, in the social web, user profile information may be augmented in various ways. For example, tag recommendation requires information about the resource to which the tag will be applied.

Because of the multi-dimensional nature of the supporting data, the social web supports a wide variety of recommendation semantics. For example, in recommending items, the system can present items personalized in terms of past viewing behavior, in terms of the behavior of friends, or in terms of a broader group of peers with similar interests. These would be presented to and interpreted by the user in different ways.

Adomavicius and Tuzhilin (Adomavicius and Tuzhilin 2005) defined the recommendation problem as that of predicting items with the highest utility for a given user. This can be achieved by a function that computes the utility for a user-item pair $u(c,i)$ where c is a user and i is an item. In the social web, we extend this formulation in three ways. first, by augmenting the input to include requirements related the recommendation. For example, when recommending tags in a social tagging system, we are interested in the utility of a tag relative to a user-resource combination: "What tag would this user want to assign to this resource?" So, the utility needs to be a function of the whole triple: user, tag, and item. Second, we note that in the social web, recommenders are not restricted to a user/item dichotomy, but may recommend users, tags, reviews, and many other things, so the object of recommendation must be broadened to include anything in the system. Finally, we see that in the social web a variety of types of utility may be relevant, as in the case of user recommendation, which can include friends, business contacts, or professional experts. So, a recommender system for a social web application may need a family of functions $u_x(c,r,o)$ where r is some set of requirements on top of the user profile, o can be any object contained in the application, and x stands for different types of utility that the recommendations need to satisfy.

Techniques for Recommendation

[1] Interestingly, tag recommendation is somewhat controversial among those concentrating on the folksonomy aspect of tagging, since an unbiased consensus requires that valuations be made independently between users and the recommender defeats user independence.

The ECML PKDD Discovery Challenge 2009

In 2009, the Knowledge and Data Engineering group at the University of Kassel organized the second ECML PKDD Discovery Challenge. It focused on the task of recommending tags in the social bookmark and publication sharing system BibSonomy (www.bibsonomy.org). Three tag recommendation tasks were formulated, based on the experience gained in the 2008 challenge.

- The first task was a cold-start task which included users previously unknown to the system, new items, and sometimes new items and new users together.
- The second task concentrated on a more conventional scenario where something is known about the user and item.
- The third challenge was online. Participants were invited to connect their recommendation engines to BibSonomy for live provision of results to its users. This challenge introduced the issue of response time in addition to recommendation quality.

In order to enable the competition, the tag recommendation feature of BibSonomy was adapted to use a multiplexer, which distributed requests to all connected recommendation engines. When a user edited a post, all recommenders got the same requests and had to provide an answer within one second. The request and all answers were stored in a database for evaluation purposes. One recommender was randomly selected and the result was presented to the user.
The system did not only store the post and the chosen tags but allowed for monitoring the click behavior of the user. If the user was not satisfied with the recommended tags he could ask for another set.
The results of the challenge were very encouraging. Most of the recommenders delivered their recommendations in time even if they were connected via the Internet. But some systems that produced high-quality recommendations in offline challenges struggled to meet the real time conditions of the online setting.
In all, 21 research groups participated in the challenge and over 150 users downloaded the datasets. For the online task, ten participants from seven countries came up with working solutions. Thirteen recommendation components were fielded for a five week period. The online competition was a unique and exiting experience for the participants and for BibSonomy users, and the winning recommendation solution is still in place providing recommendations in the system.

Eisterlehner, F.; Hotho, A.; and Jäschke, R., eds. 2009. *ECML PKDD Discovery Challenge 2009 (DC09)*, volume 497 of *CEUR-WS.org*.

Jäschke, R. 2011. *Formal Concept Analysis and Tag Recommendations in Collaborative Tagging Systems*, volume 332 of *Dissertationen zur Künstlichen Intelligenz*. Heidelberg, Germany: Akademische Verlagsgesellschaft AKA.

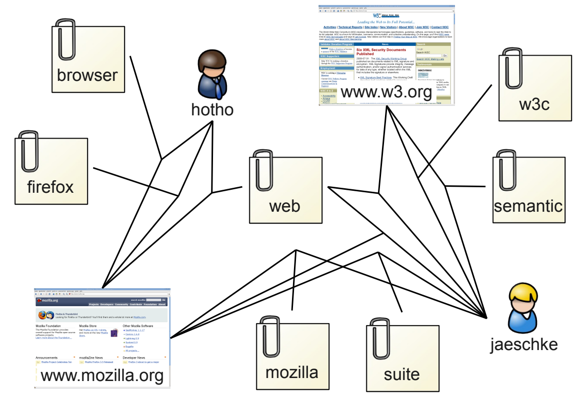

Figure 2: Excerpt of a folksonomy hypergraph from a social bookmarking system. Each assignment of a tag to an item that a user performs is represented by a hyperedge in this tripartite graph of users, tags and items. The example shows three posts made by two users.

Given the diversity of avenues and needs for recommendation in the social web, it is not surprising that many different techniques have been employed. The kind of recommendation, the motivation of the user, the type of underlying item, the form of valuation and the mode of user interaction all impact the quality of the data as well as the relative performance of the recommender. Research has shown that, in spite of the complexity, sparsity and inherent noisiness of many of the dimensions of social web data, recommenders generally do better when they take advantages of more of these dimensions rather than fewer. In this section, we look at three different approaches to recommendation in social tagging systems, each of which has a different approach to integrating the multiple dimensions of the data.

Graphs

Since the multi-dimensionality of the social web is reflected in the graph it forms, one can apply graph-based algorithms for recommendation to directly exploit the relationships between the entities, e. g., between users, items, and valuations (cf. Fig. 2). The break-through in web search at the end of the 90's was founded on a graph-based method: the PageRank algorithm (Brin and Page 1998). PageRank reflects the idea that a web page is important if there are many pages linking to it, and if those pages are important themselves. It is natural to apply similar methods for recommendation in the social web. The key idea of the *FolkRank* algorithm (Benz et al. 2010) is that an item which is tagged with important tags by important users becomes important itself. The same holds, symmetrically, for tags and users. We have thus a graph of vertices which are mutually reinforcing by spreading their weights.

Because of the different nature of folksonomies compared to the web graph (undirected triadic hyperedges instead of directed binary edges), PageRank cannot be applied directly on folksonomies. This problem is overcome in two steps. First, we transform the hypergraph into an undirected graph. Then we apply a differential ranking approach that deals with the skewed structure of the network and the undirectedness of folksonomies, and which allows for topic-specific rankings.

Folksonomy-Adapted PageRank First we convert the folksonomy hypergraph $\mathbb{F} = (U, T, I, Y)$ into an *un*directed tri-partite graph $G_{\mathbb{F}} = (V, E)$. The set V of nodes of the graph consists of the disjoint union of the sets of tags, users and items (i. e., $V = U \dot{\cup} T \dot{\cup} I$). All co-occurrences of tags and users, users and items, tags and items become edges between the respective nodes. I. e., each triple (u, t, i) in Y gives rise to the three undirected edges $\{u, t\}$, $\{u, i\}$, and $\{t, i\}$ in E.

Like PageRank, we employ the random surfer model, that is based on the idea that an idealized random web surfer normally follows links (e. g., from an item page to a tag or a user page), but from time to time jumps to a new node without following a link. This results in the following definition.

The rank of the vertices of the graph is computed (like in PageRank) with the weight spreading computation

$$\vec{w}_{\tau+1} \leftarrow dA^{\mathrm{T}}\vec{w}_{\tau} + (1-d)\vec{p} \ , \qquad (1)$$

where $\vec{w}$ is a weight vector with one entry for each node in V, A is the row-stochastic version of the adjacency matrix[2] of the graph $G_{\mathbb{F}}$ defined above, $\vec{p}$ is the random surfer vector – which we use as preference vector in our setting, and $d \in [0, 1]$ is determining the strength of the influence of $\vec{p}$. By normalization of the vector $\vec{p}$, we enforce the equality $||\vec{w}||_1 = ||\vec{p}||_1$. This[3] ensures that the weight in the system will remain constant. The rank of each node is its value in the limit $\vec{w} := \lim_{\tau\to\infty} \vec{w}_{\tau}$ of the iteration process.

For a global ranking, one will choose $\vec{p} = \mathbf{1}$, i. e., the vector composed by 1's. In order to generate recommendations, however, $\vec{p}$ can be tuned by giving a higher weight to certain nodes. For a tag recommendation we increase the weight for the user node u and the item node i. The recommended tags are then the top tag nodes in the ranking.

As the graph $G_{\mathbb{F}}$ is undirected, most of the weight that went through an edge at step τ will flow back at $\tau + 1$. The results are thus rather similar (but not identical, due to the random surfer) to a ranking that is simply based on edge degrees. The experiments we performed showed that the topic-specific rankings are biased by the global graph structure. As a consequence, we developed the following differential approach.

FolkRank – Topic-Specific Ranking The undirectedness of the graph $G_{\mathbb{F}}$ makes it very difficult for other nodes than those with high edge degree to become highly ranked, no matter what the preference vector is. This problem is solved by the *differential* approach in FolkRank, which computes a topic-specific ranking of the elements in a folksonomy. In our case, the topic is determined by the user/item pair (u, i) for which we intend to compute the tag recommendation.

[2] $a_{ij} := \frac{1}{\text{degree}(i)}$ if $\{i, j\} \in E$ and 0 else

[3] . . . together with the condition that there are no rank sinks – which holds trivially in the undirected graph $G_{\mathbb{F}}$.

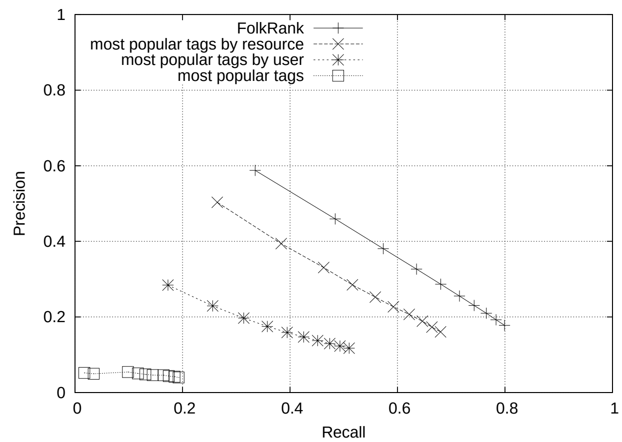

Figure 3: Recall and precision for an increasing number of recommended tags on a 2005 Delicious post-core at level 10.

1. Let $\vec{w}^{(0)}$ be the fixed point from Equation (1) with $\vec{p} = \mathbf{1}$.
2. Let $\vec{w}^{(1)}$ be the fixed point from Equation (1) with $\vec{p} = \mathbf{1}$, but $\vec{p}[u] = 1 + |U|$ and $\vec{p}[i] = 1 + |I|$.
3. $\vec{w} := \vec{w}^{(1)} - \vec{w}^{(0)}$ is the final weight vector.

Thus, we compute the winners and losers of the mutual reinforcement of nodes when a user/item pair is given, compared to the baseline without a preference vector. We call the resulting weight $\vec{w}[x]$ of an element x of the folksonomy the *FolkRank* of x.

Multi-mode Recommendations For generating tag recommendations for a given user/item pair (u, i), we compute the ranking as described and then restrict the result set to the top tag nodes. Similarly, one can compute recommendations for users (or items) by giving preference to certain users (or items). Since FolkRank computes a ranking on all three dimensions of the folksonomy, this produces the most relevant tags, users, and items for the given users (or items).

Remarks on Complexity One iteration of the adapted PageRank requires the computation of $dA\vec{w} + (d-1)\vec{p}$, with $A \in \mathbb{R}^{s\times s}$ where $s := |U| + |T| + |R|$. If τ marks the number of iterations, the complexity would therefore be $(s^2 + s)\tau \in \mathcal{O}(s^2\tau)$. However, since A is sparse, it is more efficient to go linearly over all tag assignments in Y to compute the product $A\vec{w}$. After rank computation we have to sort the weights of the tags to collect the top tags.

Results We evaluated the performance of FolkRank against other baseline methods on a dataset from Delicious (Jäschke et al. 2008). The precision-recall plot in Figure 3 shows how the recall increases, when more tags of the recommendation are used. Simultaneously, the precision drops. The plot reveals the quality of the recommendations given by FolkRank compared to other baseline approaches. The top ten tags given by FolkRank contained on average 80 % of the tags the users decided to attach to the selected item. For its top recommendations, FolkRank reaches a precision of 58.7 %.

Factorization

The multi-dimensional data of the social web can also be represented in the form of a tensor that can be factorized to exploit the latent semantic structure between users, items, and valuations. Rendle et al. introduced a model to learn the factorization of the folksonomy tensor for tag recommendation (Rendle et al. 2009). This new model focuses on a specific error function that is better suitable to predict tags in a personalized way.

Factorization Model Given a subset S of Y that represents the training data, the goal is to learn a predictor $\hat{Y}$ that predicts the tag assignments from Y that are not known in S. Therefore, the tensor Y is estimated by the three matrices $\hat{U} \in \mathbb{R}^{|U| \times k_U}$, $\hat{I} \in \mathbb{R}^{|I| \times k_I}$, $\hat{T} \in \mathbb{R}^{|T| \times k_T}$ and the core tensor $\hat{C} \in \mathbb{R}^{k_U \times k_I \times k_T}$. The factorization of the predictor $\hat{Y}$ can then be expressed as follows:

$$\hat{Y} := \hat{C} \times_u \hat{U} \times_i \hat{I} \times_t \hat{T} \tag{2}$$

The low-rank feature matrices represent the corresponding users, items, and tags resp. by a small number of latent dimensions k_U, k_I, and k_T. The core tensor $\hat{C}$ contains the connections between the latent factors. After the parameters $\hat{C}, \hat{U}, \hat{I}, \hat{T}$ have been learned,

$$\hat{y}_{u,i,t} = \sum_{\tilde{u}} \sum_{\tilde{i}} \sum_{\tilde{t}} \hat{c}_{\tilde{u},\tilde{i},\tilde{t}} \cdot \hat{u}_{u,\tilde{u}} \cdot \hat{i}_{i,\tilde{i}} \cdot \hat{t}_{t,\tilde{t}} \tag{3}$$

predicts how well the tag t fits for the given user/item pair u, i.

Learning the Model The model parameters $\hat{C}, \hat{U}, \hat{I}, \hat{T}$ can be learned by optimizing some quality criterion that compares $\hat{Y}$ with Y. Symeonidis et al. (Symeonidis, Nanopoulos, and Manolopoulos 2008) proposed to factorize Y through minimizing the element-wise loss on the elements of $\hat{Y}$ by optimizing the square loss, i. e.,

$$\underset{\hat{C},\hat{U},\hat{I},\hat{T}}{\operatorname{argmin}} \sum_{(u,i,t) \in U \times I \times T} (\hat{y}_{u,i,t} - y_{u,i,t})^2$$

This resembles higher order SVD (HOSVD), the multi-dimensional analog of singular value decomposition (SVD) for tensors. See (Kolda and Bader 2009) for a recent survey.

Rendle et al. (Rendle et al. 2009), on the other hand, propose *ranking with tensor factorization*, a method for learning an optimal factorization of Y for the specific problem of tag recommendations. Therefore, the observed tag assignments for a post $(u, i) \in P_S$ are divided into positive ($T^+_{u,i}$), negative ($T^-_{u,i}$), and missing values. Only the positive and negative values are used in the optimization:

$$\underset{\hat{C},\hat{U},\hat{I},\hat{T}}{\operatorname{argmax}} \sum_{(u,i) \in P_S} \frac{1}{|T^+_{u,i}||T^-_{u,i}|} \sum_{t^+ \in T^+_{u,i}} \sum_{t^- \in T^-_{u,i}} \frac{1}{H(u,i,t^+,t^-)}$$

with $H(u,i,t^+,t^-) = 1 + e^{\hat{y}_{u,i,t^-} - \hat{y}_{u,i,t^+}}$. The optimization is performed using gradient descent with a stochastic update approach.

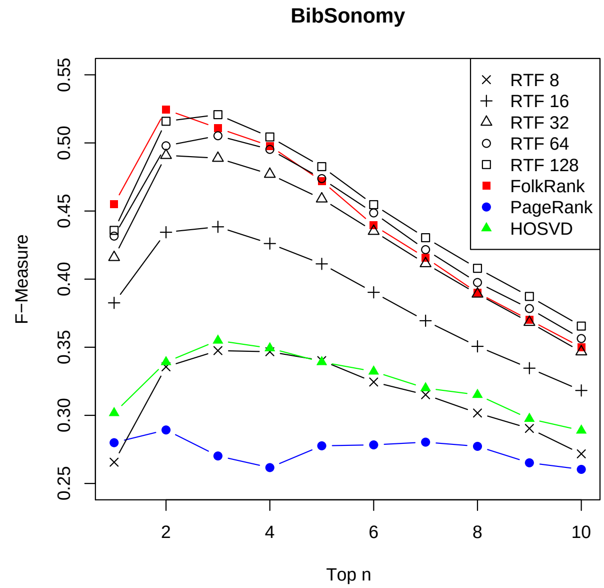

Figure 4: The F1-Measure for an increasing number of recommended tags on a BibSonomy dataset.

Results Figure 4 shows published results for the ranking with tensor factorization (RTF) technique on a dataset from BibSonomy. As the graph shows, as more latent factors are included in the RTF model (RTF 8-128), F1 (the harmonic mean of recall and precision) increases and at 128 factors, the performance is comparable to FolkRank on this tag recommendation task. The importance of optimizing only over the positive and negative observations is demonstrated by the performance of the HOSVD method, which is significantly poorer. The advantage of Rendle's approach is that one can control the tradeoff between speed and quality by selecting an appropriate number of dimensions.

Hybrids

Hybrid recommenders integrate the results of several component recommenders into a single recommendation set (Burke 2002). They have been shown effective in traditional e-commerce settings, with similar performance to other integrative solutions. In the context of the social web, hybrids have been applied most notably to social annotation systems, especially tag and item recommendation. (Gemmell et al. 2010a; Gemmell et al. 2010b)

Algorithm As one example, consider the linear weighted hybrid described in (Gemmell et al. 2010b). A schematic of this design is shown in Figure 5: as shown, a number of different recommendation components work together to assign a recommendation score combining their individual values with a linear weight or α. The weights are learned through random-restart hill-climbing.

The recommendation task in this case is to find interesting items within a social tagging system, knowing only the user and her prior tagging behavior. The aim of the hybrid is

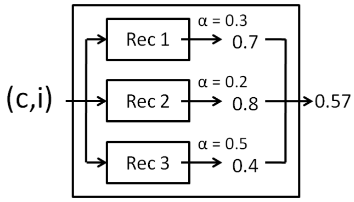

Figure 5: The linear weighted hybrid takes a user and item as input and passes it along to each component. The components individually produce relevance scores which the hybrid aggregates into a final result based on the α values.

to achieve good performance from components that are individually simple but leverage different aspects of the data. For example, we can build a two-dimensional recommendation algorithm by ignoring the tag dimension of the data and looking just for users who have tagged similar items. Or we can ignore the items and look only for neighbors with similar tag usage. Even though each of these algorithms independently might not be terribly effective, results show that combined in a linear hybrid they can achieve performance comparable to more complex algorithms.

Results Figure 6 presents an example of results (Gemmell et al. 2010b) using data from CiteULike[4], a social tagging and publication sharing site, similar to BibSonomy. In this experiment, the hybrid combines six constituent components:

Pop ignores the user and merely returns a score based on the popularity of the item.

TagSim treats both the user and item as a vector of tags and computes the cosine similarity between the two.

KNNur and **KNNut** employ user-based collaborative filtering modeling users as either items or as tags.

KNNru and **KNNrt** model items either as users or as tags in item-based collaborative filtering components.

The success of the hybrid here was achieved by leveraging all of the components, not just the strongest performing individuals, into an integrative model that exploits complementary dimensions of the data. This can be seen by examining the weights that were learned in the course of building the hybrid, shown in Table 1.

Pop	*TagSim*	*KNNur*	*KNNut*	*KNNru*	*KNNrt*
0.217	0.184	0.270	0.025	0.162	0.142

Table 1: Contribution of the individual components.

All of the components contributed to the hybrid, regardless of how well they performed alone. *KNNur* for example had an α of 0.270, the largest individual contribution. *KNNut* in contrast had an α of 0.025, perhaps because it offered little information that the other user-based collaborative filtering method did not already contribute. The remaining recommenders offered significant contributions, underscoring the need to leverage multiple dimensions of the data. Another interesting case is that of the popularity based recommender. Alone it was the worst performing recommender. As a component it had a relatively large α of 0.217, meaning it accounted for more the 20 % of the hybrid's final score.

[4] http://www.citeulike.org/

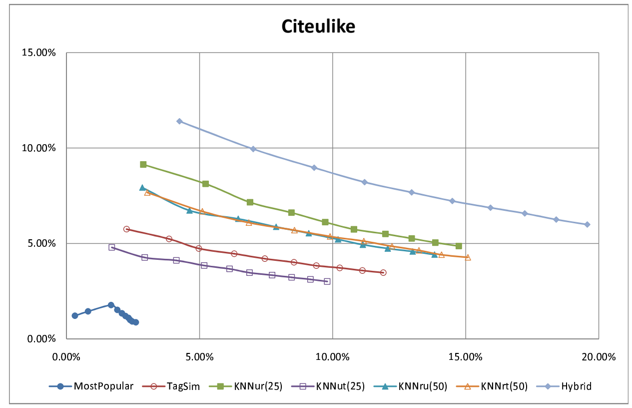

Figure 6: The recall (x-axis) and precision (y-axis) plotted for item recommendations sets of size one through ten.

Discussion The recommendation hybrid is a simple, extensible, and flexible approach to the problem of integrating the multiple dimensions of social web data, and its accuracy in tag and item recommendation is comparable to more computationally-intensive algorithms. However, the hybrid approach is not without its own challenges. The first question is which type of hybrid to use. The weighted hybrid combine scores from its components to produce a composite score. This makes sense if the components have relatively uniform performance across the data. Another possibility is the switching hybrid that chooses among several possible components in creating recommendations. If the components are relatively specialized, this may be an appropriate approach. Cascading recommenders use the results from one component to refine the output of another; this is logical when recommenders differ in accuracy while remaining complementary. The choice of which hybrid to use can have a dramatic impact on the recommendation accuracy.

Selecting a hybrid strategy often requires more decisions. If a weighted approach is selected then the question becomes how to weight the components. If a switching hybrid is used then the question is how and when to switch between components. Often this requires the additional training of a model.

In addition to selecting the type of hybrid, the designer must also chose the component recommenders themselves. The best overall performance is often achieved by selecting components that complement one another. As the results above show, even components with weak individual performance may make significant contributions to overall effectiveness when assembled in a hybrid.

Conclusion

The social web is an important emerging area in recommender systems research. Compared to two-dimensional recommendations problems typically found in e-commerce settings, the social web presents unique challenges. In particular:

- The diversity of the items to which recommendation may be applied,
- The wide range of recommendation semantics that prevail depending on the nature of the application, and
- The complexity and multi-dimensional nature of the social web data to which recommendation algorithms are applied.

For these reasons, the choice of recommendation approach is likely to be highly dependent on the specific recommendation task and domain. Social tagging is the best-studied social web recommendation application, and in this article, we have illustrated three different approaches to this problem.

References

[Adomavicius and Tuzhilin 2005] Adomavicius, G., and Tuzhilin, A. 2005. 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.

[Benz et al. 2010] Benz, D.; Hotho, A.; Jäschke, R.; Krause, B.; Mitzlaff, F.; Schmitz, C.; and Stumme, G. 2010. The social bookmark and publication management system bibsonomy. *The VLDB Journal* 19(6):849–875.

[Brin and Page 1998] Brin, S., and Page, L. 1998. The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks and ISDN Systems* 30(1-7):107–117.

[Burke 2002] Burke, R. 2002. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction* 12(4):331–370.

[Gemmell et al. 2010a] Gemmell, J.; Schimoler, T.; Mobasher, B.; and Burke, R. 2010a. Hybrid tag recommendation for social annotation systems. In *19th ACM Int. Conf. on Information and Knowledge Management*. New York, NY, USA: ACM.

[Gemmell et al. 2010b] Gemmell, J.; Schimoler, T.; Mobasher, B.; and Burke, R. 2010b. Resource recommendation in collaborative tagging applications. In *11th Int. Conf. E-Commerce and Web Technologies*, 1–10. Springer Verlag.

[Jäschke et al. 2008] Jäschke, R.; Marinho, L.; Hotho, A.; Schmidt-Thieme, L.; and Stumme, G. 2008. Tag recommendations in social bookmarking systems. *AI Communications* 21(4):231–247.

[Kolda and Bader 2009] Kolda, T. G., and Bader, B. W. 2009. Tensor decompositions and applications. *SIAM Review* 51(3):455–500.

[Rendle et al. 2009] Rendle, S.; Marinho, L. B.; Nanopoulos, A.; and Schmidt-Thieme, L. 2009. Learning optimal ranking with tensor factorization for tag recommendation. In *KDD '09: Proceedings of the 15th ACM SIGKDD Int. Conf. on Knowledge discovery and data mining*, 727–736. New York, NY, USA: ACM.

[Symeonidis, Nanopoulos, and Manolopoulos 2008] Symeonidis, P.; Nanopoulos, A.; and Manolopoulos, Y. 2008. Tag recommendations based on tensor dimensionality reduction. In *RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems*, 43–50. New York, NY, USA: ACM.