

Social Media Recommendation based on People and Tags

Ido Guy, Naama Zwerdling, Inbal Ronen, David Carmel, Erel Uziel

IBM Research Lab

Haifa 31905, Israel

{ido,naamaz,inbal,carmel,erelu}@il.ibm.com

ABSTRACT

We study personalized item recommendation within an enterprise social media application suite that includes blogs, bookmarks, communities, wikis, and shared files. Recommendations are based on two of the core elements of social media—people and tags. Relationship information among people, tags, and items, is collected and aggregated across different sources within the enterprise. Based on these aggregated relationships, the system recommends items related to people and tags that are related to the user. Each recommended item is accompanied by an explanation that includes the people and tags that led to its recommendation, as well as their relationships with the user and the item. We evaluated our recommender system through an extensive user study. Results show a significantly better interest ratio for the tag-based recommender than for the people-based recommender, and an even better performance for a combined recommender. Tags applied on the user by other people are found to be highly effective in representing that user's topics of interest.

Categories and Subject Descriptors: H.3.3

[Information Search and Retrieval]: *information filtering*

General Terms: Algorithms, Experimentation

Keywords: Personalization, Recommender Systems, Social Media, Social Networks, Social Software, Collaborative Tagging

1. INTRODUCTION

Social media has been enjoying a great deal of success in recent years, with millions of users visiting sites like Facebook for social networking; Wordpress for blogging; Twitter for micro-blogging; Flickr and YouTube for photo and video sharing, respectively; Digg for social news reading; and Delicious for social bookmarking. These social media sites rely principally on their users to create and contribute content; to annotate others' content with tags, ratings, and comments; to form online relationships; and to join online communities.

As social media sites continue to proliferate, and their volumes of content keep growing, users are having more difficulty choosing sites in which to become actively involved. Furthermore, users are "flooded" with information from feed readers, news alert systems,

and many other resources. Easy access to so much information along with difficulty in judging the validity of so much content can lead to *information overload*, i.e., having more information available than a user can readily assimilate. Social media sites are increasingly challenged to attract new users and retain existing ones, due to these same factors.

One way site address these issues is by providing users with personalized recommendations. As in traditional taste-related domains or e-commerce (movies, books, hotels), the goal of a personalized recommender system is to adapt the content based on characteristics of the individual users. Social media and personalized recommender systems can mutually benefit from one another: on the one hand, social media introduces new types of public data and metadata, such as tags, ratings, comments, and explicit people relationships, which can be utilized to enhance recommendations; on the other hand, recommender technologies can play a key role in the success of social media applications and the social web as a whole, ensuring that each user is presented with the most attractive and relevant content, on a personal level.

In recent years, quite a few personalized recommendation services for social media have emerged. For instance, StumbleUpon¹ is a personalized recommender engine that suggests web pages based on a user's past ratings, ratings by friends, ratings by users with similar interests, and topics of interest selected by the user from a list of nearly 500 subjects. More recently, some of the leading social media sites have also added personalized recommendation features: video-sharing site YouTube has launched a personalized homepage that includes recommendations based on past views and favorites. This feature is reported to have led to an increase in the number of users visiting the homepage, the frequency of visits, and the number of subscriptions users make over time [25]. Social news aggregator service Digg has added a personalized recommender engine for presenting stories presumed to be most interesting to a user, based on preferences of similar users [24].

Following the proliferation of social media sites on the web, analogous sites have emerged within organizations, gaining popularity as well [8]. Similarly to their counterparts on the web, enterprise social media sites also face challenges stemming from a continuously growing number of applications and the expanding volumes of information within them [8,11].

1.1 Contribution

In this work, we study personalized recommendation of social media items within an enterprise social software application suite, Lotus Connections (LC) [18]. LC consists of various types of social media applications, including social bookmarking, file

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGIR '10, July 19–23, 2010, Geneva, Switzerland.

Copyright 2010 ACM 978-1-60558-896-4/10/07...\$10.00.

¹ www.stumbleupon.com

sharing, blogging, communities, and wikis. Our recommender suggests items across the different applications based on two of the main characteristics of social media—people and tags.

In a previous work, we studied the recommendation of social media items based purely on related people [17]. We showed that items that are strongly related to people in a user’s social network are likely to interest that user. Our hypothesis in this work is that recommending items related to a user’s tags can also increase the quality of recommendation. Such a combination may be viewed as a social media variation of a traditional hybrid recommender that has been proven to be effective in taste-related domains [4].

Previous work has suggested tag-based recommendations, highlighting the value of tags as concise and accurate content descriptors that take into account human perception of the content [22,29]. User-tag relationships have been inferred through direct usage of tags or through indirect links, such as tags applied to resources rated positively by a user or those that were clicked-through by a user. In this work, we only use information that is already publicly available and that does not require any explicit input, such as rating. We do not use any private information, such as click-through rates or query logs. We evaluate three methods to extract user-tag relationships based on public information: (1) direct usage of tags across the different LC applications (“used tags”); (2) indirect link between a user and a tag through an item, e.g., tags related to documents that are related to the user (“indirect tags”); and (3) tags applied to the user by others, within a people-tagging feature that allows users to tag one another [9] (“incoming tags”). To the best of our knowledge, our study is the first to suggest using incoming people tags to recommend content.

Our recommender engine is based on the social aggregation system SaND [5,27], which aggregates relationships among people, items, and tags, across the different LC components. SaND is used to extract, for each user, weighted lists of related people and related tags that constitute the user’s personal profile. In addition, SaND provides weighted lists of items related to given people and/or tags. Ultimately, the system recommends to the user items that are related to people and tags within his personal profile. For each recommended item, two-level explanations illustrate why the item is recommended. On the first level, the related people and/or tags that yielded the recommended items are presented. On the second level, by hovering over the name of a specific person or a tag, the user may see its relationship to the recommended item and to himself as inferred by SaND.

Our approach has several advantages: (1) users are not required to provide explicit input to the system, e.g., by rating a set of items (we infer both their social relationships and topics of interest from other online information); (2) coping with the cold start problem of new users [28], as SaND allows aggregation of data which is external to LC (see [11]); (3) transparency [31]—intuitive explanations can be provided based on public tags and social relations; (4) performance—our recommendations are based on the rich aggregated index and do not require clustering or other computationally-intensive methods; and (5) generality—both people and tags can be used to recommend virtually any type of item, including music, photos, and videos.

While the SaND infrastructure has been used before for providing people-based recommendations, in this work we describe how it can be exploited to provide effective tag-based recommendations

as well. Furthermore, we present a novel approach for a hybrid recommender based on people and tags that leverages the unified modeling of relationships among people, tags, and resources. Another benefit of this approach is a uniform presentation of “hybrid explanations” based on both people and tags.

1.2 Evaluation

Our evaluation aims at comparing five types of recommenders: a people-based recommender (PBR); a tags-based recommender (TBR); two types of a hybrid recommender (PTBR): a combination of people **or** tags (or-PTBR), and a combination of people **and** tags (and-PTBR, suggesting only items related to both people and tags); and a popularity-based recommender (POPBR), as a benchmark. To the best of our knowledge, this is the first comprehensive study to compare people-based recommenders with tag-based recommenders and their hybridizations.

Our evaluation involves the following elements: (1) an offline comparison of the recommended items yielded by the five recommenders over 1,410 LC users, to examine the diversity across the recommenders, and in particular to compare the items stemming from related people with the items stemming from related tags; (2) a user survey with 65 participants who were asked to evaluate tags as indicators of topics of interest, based on four different methods: indirect tags, used tags, incoming tags, and a combination of both used and incoming tags; (3) the main element of our evaluation is a survey of over 400 LC users, who were randomly divided into five groups, receiving recommendations based on the five recommenders. All groups received recommendations in two phases—without explanations and with explanations. Participants were asked to provide feedback on their interest in the recommended items.

Our primary results show that the combination of incoming tags and used tags is the most effective in representing a user’s topics of interest, with users rating nearly 70% of the topics as very interesting. Recommendations based on a TBR, with a tag profile that combines incoming and used tags, are rated significantly more interesting than the most effective PBR studied in our previous work. Recommended items are shown to be highly different between the PBR and the TBR, with less than 2% overlap. A hybrid PTBR recommender including explanations improves the results slightly further, leading to an over 70:30 ratio between interesting and non-interesting items. It also presents other potential benefits over a TBR, such as a lower percentage of already known items and higher diversity of item types.

In the next section, we discuss how existing work relates to our research. We then present our recommender system, followed by a detailed description of our experiments and their results. We conclude by discussing our findings and suggesting future work.

2. RELATED WORK

There are two prevalent approaches for building recommender systems: content-based (CB) [26] and collaborative filtering (CF) [13]. The CB approach is based on recommending items that are similar to those in which the user has shown interest in the past. The CF approach, on the other hand, recommends items to the user based on other individuals who are found to have similar preferences or tastes. Traditionally, both CB and CF systems have been based on explicit input from the user, usually provided by rating a set of items. To avoid this extra burden on the user,

leveraging implicit interest indicators [6], such as purchase history, views, clicks, or queries, has recently become more popular in recommender systems.

With the current prosperity of social media in general, and of social network sites (SNSs) in particular, several studies have suggested incorporating direct social relationships in CF systems. ReferralWeb [19] was one of the first systems to suggest the combination of direct social relations and CF to enhance searching for documents and people. Several studies suggest incorporating explicit social network information in CF systems to improve the quality of recommendation in domains such as movies and books (e.g., [3,12,30]), music [20], clubs [14], and news stories [21]. In this work, we infer social relationships from many different data sources, such as an enterprise SNS, a wiki system, and an organizational chart. Previous work has shown the value of aggregating social network information in yielding a richer and more accurate social graph [15].

On the other hand, as tagging has emerged as a popular way to let users annotate social media content, several works propose using tags as content descriptors for CB systems. Li et al. [22] analyze data from the social bookmarking site Delicious and find a high similarity between the tag vector of a URL and its keyword vector, as extracted from the corresponding web page. Firan et al. [10] study personalized recommendation of tracks within the popular music portal Last.Fm, and show that tag-based profiles can produce better recommendations than conventional ones based on track usage. Vatturi et al. [32] study personalized bookmark recommendation using a CB approach that leverages tags, assuming that users would be interested in pages annotated with tags similar to ones they have already used. Sen et al. [29] introduce Tagommenders—recommender algorithms that extend existing CB techniques by making use of tags. Their evaluation is based on the MovieLens system, and findings indicate that tag-based algorithms generate better recommendation rankings than state-of-the-art CF-based algorithms. The value in generating intuitive explanations through tags is highlighted in another MovieLens study by the same authors [33]. Our own tag-based approach is based on aggregating tags across various social media systems and considering both tags used by the user as well as tags with which the user has been tagged.

In this paper, we use the combination of related people and related tags to recommend social media items. Our system can be viewed as a variation of a hybrid CF-CB recommender system, in which related people and tags are used analogously to traditional CF and CB systems, respectively. Some research suggests combining traditional CF and CB systems, mostly in taste-related domains (see [4] for a summary). In particular, several studies point to the value of hybridizing CF and CB over each of the pure methods on its own. For example, Fab [2], a hybrid recommender system for web pages, is one of the first systems that combined CB and CF, suggesting that such a combination may eliminate many of the weaknesses found in each approach when individually applied. Claypool et al. [7] present a new filtering approach that combines the “coverage and speed” of CB filters with the “depth” of CF, and provides personalized filtering of an online newspaper. Melville et al. [23] present a hybrid recommender approach—Content-Boosted Collaborative Filtering (CBCF), which uses a CB predictor to enhance existing user data, and then provides personalized suggestions through CF.

Evaluation is based on a movie rating dataset and indicates that CBCF performs better than pure CB or pure CF. The hybrid recommender presented in this work is based on implicit interest indicators and does not require explicit ratings by users, as most of the previous work. The unique hybridization algorithm is based on a unified index [1], which allows integrated retrieval of recommended items based on both people and tags.

3. RECOMMENDER SYSTEM

3.1 Social Media Platform

Our research platform for personal recommendation is Lotus Connections (LC) [18]—a social software application suite for organizations. It includes seven social media applications: profiles (of all employees), activities, bookmarks, blogs, communities, files, and wikis. We focus on recommending items of the last five applications, disregarding the first two, since profiles pose a different challenge regarding people recommendation [16], and an activity is generally restricted to a limited number of users. In our work, recommended items may originate from one of the following five applications, which are part of LC’s deployment within our organization: (1) social bookmarking application, which allows users to store and tag their favorite web pages. It includes 900K bookmarks with 2M tags by 21K users; (2) blogging service that contains 7.5K public blogs, 130K entries, 350K tags and 17K users; (3) online community system that contains 6K public communities, each with shared resources (such as feeds and discussion forums), with a total of 174K members and 19.5K tags; (4) system for file sharing with 15K public files (presentations, photos, articles, etc.), 24K tags, and 8K users; and (5) wiki system with 3K public wikis including 20K pages edited by 5K users, and with 10K tags.

3.2 Relationship Aggregation

SaND [5,27] is an aggregation system that models relationships among people, items, and tags, through data collected across the enterprise, and in particular across all LC applications. SaND aggregates any kind of relationships between its three core entities—people, items, and tags. The implementation of SaND is based on a unified approach [1], in which all entities are searchable and retrievable. As part of its analysis, SaND builds an entity-entity relationship matrix that maps a given entity to all related entities, weighted according to their respective relationship strengths. The entity-entity relationship strength is composed of two types of relations:

- *Direct Relations*: Figure 1 shows all direct relations among entities that are modeled by SaND. Particularly, a user is directly related to: (1) another person: as a friend, as a tagger of or tagged by that person, or through the organizational chart (direct manager or employee); (2) an item (e.g., a shared file or a community): as an author, a commenter, a tagger, or a member; or (3) a tag: when used by the user or applied on the user by others. In addition, an item is directly related to a tag if it has been tagged with it. SaND does not currently model any direct tag-tag and item-item relations.
- *Indirect Relations*: Two entities are indirectly related if both are directly related to another common entity. For example, two users are indirectly related if both are related to the same user, e.g., if both have the same manager or friend, or if both have tagged or were tagged by the same person.




Figure 1. Direct entity-entity relations in SaND.

3.3 User Profile

The user profile, $P(u)$, is given as an input to the recommender engine once the user u logs into the system. The profile is used to personalize the recommended items for u . It consists of 30 related people, $N(u)$, and 30 related tags, $T(u)$, retrieved through SaND, as explained in the paragraphs below.

The set of people related to the user is extracted by considering both direct and indirect people-people relations, scoring them, and aggregating them into a single person-person relationship strength, in the same way as was performed in previous studies ([16,17]). In principle, each direct relation adds a score of 1 to the overall relationship score, while an indirect relation adds a score in the range of (0,1), determined by various parameters, such as the number of common files or number of other wiki co-authors. More details on person-person score calculation can be found in [15,16,17].

Our previous work on purely people-based recommendation [17] distinguished between familiarity relationships (people the user knows) and similarity relationships (people whose social activity overlaps with the user’s social activity). Familiarity relationships include all direct people-people relations, as well as two types of indirect relations: co-authorship (e.g., of a file or a wiki), and having the same manager. Similarity relationships include indirect relations only, such as co-usage of the same tag, co-tagging of the same item, co-commenting on the same blog entry, or co-membership in the same community. Findings of that work have indicated that familiarity relationships are more effective in yielding interesting recommended items, yet similarity relationships are also productive and may diversify the recommended items. Based on our previous work’s conclusions, all similarity relationships are multiplied by a factor of 1/3, so that familiarity relationships are favored, yet do not completely prevail. The user’s set of related people is ultimately determined by retrieving the 30 related people who are found to have the highest relationship strength with the user, as done in [17].

To extract the user’s related tags, we consider the following user-tag relations: (1) used tags—direct relation based on tags the user has used; (2) incoming tags—direct relation based on tags applied on the user by others; and (3) indirect tags—indirect relation based on tags applied on items related to the user (note that this subsumes relation 1). We conducted a user survey to evaluate the quality of these tags as indicators for the user’s topics of interest. Results of this evaluation are used to configure SaND to return the 30 tags that are most strongly related to the user’s topics. The survey results are described in more detail in Section 4.1.

3.4 Recommendation Algorithm

Given the user profile, $P(u) = (N(u), T(u))$, we suggest items to the user that are related to people and/or tags in his profile. The recommendation score of item i for user u is determined by:

$$RS(u, i) = e^{-\alpha d(i)} \cdot [\beta \sum_{v \in N(u)} w(u, v) \cdot w(v, i) + (1 - \beta) \sum_{t \in T(u)} w(u, t) \cdot w(t, i)]$$

where $d(i)$ is the number of days since the creation date of i ; α is a decay factor (set in our experiments to 0.025, as in [17]); β is a parameter that controls the relative weight between people and tags, and is used in our experiments to evaluate different recommenders; $w(u, v)$ and $w(u, t)$ are the relationship strengths of u to user v and tag t , as given by the user profile; $w(v, i)$ and $w(t, i)$ are the relationship strengths between v and t , respectively, to item i , as determined by SaND, based on direct relations as described in Figure 1. User-item direct relation types are weighted as in previous studies [1,5,17]: authorship (0.6), membership (0.4), commenting (0.3), and tagging (0.3). Tag-item relations are weighted relative to the number of users who applied the tag on the item, normalized by the overall popularity of the tag, as in [1].

Ultimately, the recommendation score of an item, reflecting its likelihood to be recommended to the user, may increase due to the following factors: more people and/or tags within the user’s profile are related to the item; stronger relationships of these people and/or tags to the user; stronger relationships of these people and/or tags to the item; and freshness of the item. We exclude items that are found to be directly related to the user. For example, we will not recommend an item on which the user has already commented or has already tagged.

3.5 Recommender Widget

Figure 2 depicts our UI widget for item recommendations based on the algorithm described in the previous section. The user is presented with a number of items (three, in this example) that may include a mix of the five LC item types. Each item has a title that links to the original document, and a short description when available. The icon to the left of each item represents its type—the first item in Figure 2 is a blog entry, the second is a community, and the third is a wiki.

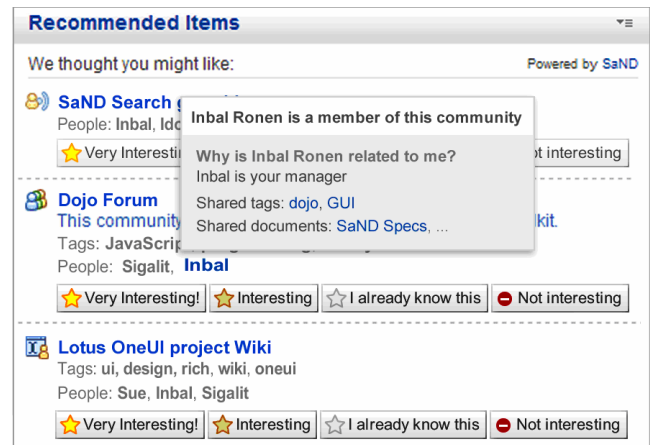


Figure 2. Item Recommendation Widget.

Each item includes a list of up to five related person names and/or up to five related tags that yielded this item’s recommendation. The related people and tags serve as a first level explanation of why the item is recommended. On the second level, when hovering over a person’s name or a tag, the user is presented with a popup detailing the relations of the person/tag to the user and to the item. In Figure 2, the popup indicates that Inbal is a member of the recommended community, and is also related to the user through several detailed direct and indirect relations. In the case of hovering over a tag, the popup indicates whether the user has used the tag, was tagged by the tag, or both.

4. EVALUATION

4.1 Tag Profile Survey

As a first step of our evaluation we set out to explore how to effectively build a user’s tag profile based on the information represented in SaND. As described in the previous section, we examine three types of user-tag relations: used tags, indirect tags, and incoming tags. While the first two types have been used in previous studies around tag-based personalization, to the best of our knowledge, this is the first study that examines incoming tags for personalized content recommendation.

Our evaluation is based on a user survey sent to 200 LC users with at least 30 used tags and 30 incoming tags. User-related topics were assumed to be represented by tags associated with the user through four types of user-tag relations: (1) used tags; (2) incoming tags; (3) indirect tags; and (4) direct tags. The last group considers both types of direct relations (used tags and incoming tags) as retrieved through SaND. We extracted the user’s four top related tags based on each of the relation types and randomized their order. Overall, we produced up to 16 tags for each of the participants, for which they were asked to indicate their level of interest, according to following three options: “Not Interested”, “Interested”, and “Highly Interested”. We sent invitations to the survey by email, and received responses from 65 users, who rated a total of 1,037 tags.

Table 1. Rating results of tags as topics of interest

%	Not Interested	Interested	Highly Interested
used	16.84	38.25	44.91
incoming	15.48	31.75	52.78
direct	7.46	22.81	69.74
indirect	35.38	45.38	19.23

Table 1 shows the rating results of the tags as topics of interest for each of the four relation types. Direct tags clearly yield the most interesting topics—nearly 70% are rated as highly interesting and only 7.5% are rated not interesting. Incoming tags are slightly more effective in representing topics of interest than used tags, while indirect tags are evidently the least effective, with only 19% rated as highly interesting. One-way ANOVA indicates that ratings across the four types are significantly different ($F(3,1068)=51.89, p<.0001$). Tukey post-hoc comparisons of the four types indicate that direct tags are rated significantly higher than the rest of the types, indirect tags are rated significantly lower than the rest, and that the difference in interest levels between incoming and used tags is not significant.

Due to these results we opted to use the direct user-tag relation for retrieving the user’s tag profile. We did not further weight

incoming vs. used tags, as the differences between them in the survey were not statistically significant. Consequently, we used SaND’s indirect relations only for retrieving the list of people related to a user (as has been shown useful by a previous study [15]).

4.2 Recommended Items Survey

4.2.1 Methodology

The main part of our evaluation is based on an extensive user survey, designed to compare the people-based recommender (PBR), the tag-based recommender (TBR), and two combinations of these two recommenders (PTBRs). Participants of the survey were asked to evaluate 16 recommended items in two randomly ordered phases (each phase included eight items): with and without explanations. Each participant was assigned to one of five groups in a round-robin order, receiving recommendations based on one of the following five recommenders: (1) PBR ($\beta=1$ in the equation in Section 3.4); (2) TBR ($\beta=0$); (3) or-PTBR—each item may be recommended due to related people, related tags, or both ($\beta=0.5$); (4) and-PTBR—each item is recommended due to at least one person and at least one tag in the user’s profile ($\beta=0.5$ with the constraint that both parts of the summation in brackets are nonzero); and (5) POPBR—popular item recommendation (as a benchmark). The popularity of items was determined based on the number of people they were directly related to in SaND, and on the items’ freshness. For explanations, we pointed out the types and numbers of the different direct relations with people as well as the last-update date. For example, an explanation for a popular item would be: “tagged by 57, commented by 12, last updated Jan. 17th, 2010”. Recommended items in each of the two phases were presented using the widget described in Figure 2, allowing to rate them as “Very Interesting”, “Interesting”, “I already know this”, or “Not Interesting”.

Our target population for the survey consisted of 1,410 LC users who were directly related to at least 30 other people, 30 tags, and 30 items. We note that this group does not represent the entire population of our organization, but rather active users of the LC system, who are the target population for our recommender system. A link to the survey with an invitation to participate was sent to each of these 1,410 individuals. In addition, we ran the five recommenders for each of these users to retrieve the top 16 items, and calculated average overlap between the items returned from the different recommenders. The average overlap across the 1,410 users between the items returned by the PBR and the TBR was 1.58%, indicating that these two recommenders return very dissimilar items. The POPBR had very low overlap with all other recommenders, ranging from 0.87% to 1.83%. Overlap between the two PTBRs was 38.6%. The or-PTBR had higher overlap with the PBR (57.3%) and the TBR (32.6%) than the and-PTBR (24.1% and 9.7%, respectively). This indicates that the or-PTBR recommends mostly items that are either recommended by the PBR or the TBR, while the and-PTBR recommends more items that are further down the list of the PBR and the TBR.

4.2.2 Results

In total, 412 participants completed our survey, originating from 31 countries and spanning the different organizational units: 32% sales, 28% software, 18% services, 11% headquarters, 4% research, 4% systems, and 3% others.

Table 2 summarizes the rating results of the survey for each of the five recommenders, with and without explanations. The rightmost column displays the *interest ratio*—the ratio between interesting (including very interesting) and non-interesting items. The best ratio is achieved by the or-PTBR with explanations. One-way ANOVA indicates that ratings across the five recommenders are significantly different ($F(4,5496)=66.823, p<.0001$). Tukey post-hoc comparisons indicate that differences between the POPBR and the other four recommenders, as well as between the PBR and the other four recommenders, are significant, while differences among the TBR, and-PTBR, and or-PTBR are not significant.

Table 2. Item rating results across the five recommenders

Rec.	Expl.	Not Int.	Int.	Very Int.	Aldr. Know	Int. Ratio
POPBR	no	52.72	30.95	11.05	5.27	0.80
	yes	58.12	29.96	7.58	4.33	0.65
PBR	no	37.70	35.25	18.03	9.02	1.41
	yes	34.94	38.35	18.32	8.39	1.62
TBR	no	26.41	36.46	23.17	13.97	2.26
	yes	26.97	35.13	21.14	16.76	2.09
or-PTBR	no	29.87	36.95	21.07	12.11	1.94
	yes	26.10	40.00	23.39	10.51	2.43
and-PTBR	no	31.41	37.50	21.28	9.70	1.87
	yes	25.90	37.52	24.02	12.56	2.38

4.2.2.1 Baseline – POPBR

We opted to use popular items as a benchmark, in order to examine whether personalized recommendations are “worth the effort”, and add substantial value over a general non-personalized recommendation of the most popular items. Results show that all types of personalized recommenders significantly outperform the popularity-based recommender. Interestingly, when accompanied by explanations, popular items are rated slightly lower, possibly as the numbers indicating an item’s popularity are not found to be a compelling justification by the participants.

4.2.2.2 PBR vs. TBR

Results for the PBR are consistent with the results from our previous work about people-based recommendation [17]. In this case, the explanations slightly increase the interest rate in recommended items, reinforcing the instant value of people-based explanations.

In this work, we suggest using tags to improve the quality of social media item recommendation and as described in Section 4.1, we leverage both used tags and incoming tags. Results for the TBR, as displayed in Table 2, reveal that tag-based recommendations significantly outperform people-based recommendations, both with and without explanations. When tags come into play, interest ratio jumps from around 1.5 to over 2. As opposed to the PBR, tag-based explanations did not instantly lead to more interest in items. In fact, items without explanations even led to a slightly higher interest ratio (2.26 vs. 2.09). This indicates that while a TBR outperforms a PBR, tag-based explanations are not as effective as people-based explanations in increasing interest in items. This may be due to the fact that related tags are already reflected, to some extent, in an item’s title or description.

Figure 3 depicts aggregated rating results across the five recommenders, regardless of whether explanations were provided or not. The “All Interesting” bar includes items rated interesting and very interesting. It can be seen that the number of items rated non-interesting drops by almost 10% when moving from the PBR to the TBR. The percentage of already known items is also notable—it is the highest for the TBR (15.5%), the lowest for the POPBR (4.8%) and the next-to-lowest for the PBR (8.7%). Tukey post-hoc comparisons indicate that the percentage of already known items in the TBR is significantly higher than in any of the other recommenders (other differences among the recommenders are not significant). These results are understandable: tags yield items that are more similar to ones the user preferred in the past and are thus likely to be less diverse and surprising than people-based items. Diversity is a well known advantage of CF recommenders over CB ones [13]. Popular items are even less expected than people-based items, however, as mentioned before, their interest ratio is significantly lower. The high diversity of the PBR relative to the TBR is also reflected in the types of items each of them yielded: 80% of the items recommended by the TBR for the survey’s participants were bookmarked web pages, while for PBR 32.6% were files, 29.3% communities, 22.3% bookmarks, 12.3% blogs, and 3.4% wikis.




Figure 3. Item rating results across the five recommenders, summing over both phases.

4.2.2.3 PTBR vs. TBR

Figure 3 indicates that both the or-PTBR and the and-PTBR produce the highest percentage of items rated interesting or very interesting—over 60%. Table 2 demonstrates that when explanations are included, both PTBRs also have the highest percentage of very interesting items, the lowest percentage of non-interesting items, and an overall interest ratio of over 70:30. However, the differences between the two PTBRs and the TBR are statistically insignificant. Moreover, when explanations are excluded, the TBR performs slightly better than both PTBRs. These differences between the two phases may be due to the effectiveness of the people-based explanations included in the PTBRs, but not in the TBR (as discussed before).

Figure 3 also shows that both PTBRs have a significantly lower percentage of already known items than the TBR, indicating that they produce less expected items while maintaining high interest ratios. Diversity of item types is also higher for the PTBRs, as compared to the TBR: only 44% of the or-PTBR and 51% of the and-PTBR items are bookmarks (compared to 80% for the TBR).

4.2.2.4 Item Type Diversity

42.5% of the recommended items in our survey (over all 412 participants) were bookmarked web pages, 26.9% were shared

files, 15.6% were communities, 10.2% were blog entries, and 4.8% were wikis. These differences may be ascribed to the fact that applications are different in various parameters, such as the level of usage within the organization, the frequency of item creation in the system (e.g., a bookmark is more frequently created than a wiki), and so on. Our recommendation algorithm does not explicitly consider the item type, and does not impose predefined item type diversity.




Figure 4. Item rating results across the five types, summing over both phases.

Figure 4 depicts the rating results across the five item types. Bookmarks have the highest interest ratio, followed by communities and files. Ratings of these three types are significantly higher than those of blogs and wikis (using ANOVA with Tukey post hoc analysis). Note that the order of items in terms of proportion in the overall recommendations, as detailed in the previous paragraph, is very similar to their order in terms of interest ratio. Our recommenders suggest more items of types that are likely to be interesting, but also maintain some level of item type diversity. The percentage of already known items per type increases in accordance with the interest ratio, indicating a trade-off between accuracy and expectedness.

5. DISCUSSION AND FUTURE WORK

The results presented in the previous section indicate that using tags for social media recommendation can be highly beneficial. The combination of directly used tags and incoming tags produces an effective tag-based user profile. A TBR that makes use of this profile yields significantly more interesting recommendations than the most effective PBR presented in a previous work [17]. In addition, the items produced by the TBR are almost completely disjoint from the items produced by the PBR (less than 2% average overlap across the top 16 items), indicating that related tags produce very different recommendations as compared to related people. Combining both related tags and people in the user profile does not significantly increase the interest in recommended items over a pure tag-based approach; however, it significantly lowers the percentage of already known items, increases the diversity of item types, and makes explanations more effective.

The higher effectiveness of the TBR over the PBR may be attributed to the fact that tags are better filters for topics of interest than are people. People related to the user may broaden the scope of recommended items (and increase diversity), yet they are also likely to add irrelevant items, as they may have interest areas that are different from those of the user.

In our previous work on PBRs [17], some of the feedback we received highlighted the need for additional filtering based on topics

(since related people often have many different topics of interest). Indeed, PTBRs are found to perform significantly better than the optimal PBR. However, this improvement does not occur over the TBR, a finding that surprised us to some extent. We expected that adding people to tags as filters would significantly improve the recommendations (similarly to traditional hybrid recommender systems), yet the improvement was small. Our findings suggest that a TBR without explanations performs well, and can be used as a starting point, or in cases that require a simple social media recommender system.

We examine two PTBRs that combine related people and related tags in different manners, and produce fairly dissimilar recommended items (less than 40% mutual overlap). Yet, the differences in their performance are very small. Future studies may examine whether other methods for combining related people and related tags in user profiles can further enhance the recommender’s performance.

In addition to the people- and tag-based recommenders, we also experimented with a non-personalized, popularity-based recommender. While the interest ratio of this recommender is significantly lower than all personalized recommenders, it has the potential to provide more unexpected recommendations, as reflected in its very low percentage of already known items. In a future work, we plan to examine whether and how a popularity recommender can be combined with the personalized recommenders, so that more unexpected items are suggested to the user, but not so often as to become an annoyance.

Integration of traditional CB methods within the recommender should also be explored and can be helpful in addressing two key issues that are acute in both TBRs and PBRs: (1) the cold start problem for new items, as these are not yet related to people or tags, and (2) language issues—items that users cannot understand might be accidentally recommended (e.g., when the tag’s language is different than the language of the content).

Our recommender engine is based on the rich relationship data aggregated and modeled by SaND. The fact that we do not apply computationally-intensive algorithms over this data allows us to compare recommenders in a more direct way, provide intuitive explanations, and maintain generality. Future research should examine whether applying such algorithms can further improve the results presented in this work.

A future study is also required to validate the results of our experiments in a non-enterprise environment, where tags are used on a larger scale, related people are mostly personal friends rather than colleagues, and multiple identities must be managed.

We also plan to examine how to maintain high interest in recommended items over time. While the evaluation in this study is mostly based on rating an initial set of recommended items, maintaining that same level of interest for users who regularly access the system is more challenging. One approach we intend to explore, which could help overcome this challenge, is based on user feedback. The approach would address how to elicit such feedback, on what levels to allow it (an item, a person, a tag, etc.), and how to adapt the recommendations accordingly.

6. CONCLUSION

In this work, we propose a novel method for recommending social media items based on both related people and related tags. An

extensive experimentation is conducted to compare people-based and tag-based recommenders as well as their hybridizations. We show that a combination of directly used tags and tags applied by others is most effective in representing the user's topics of interest. A recommender based on this tag profile yields items that are significantly more interesting to the user than the most effective people-based recommender demonstrated in a previous work [17]. Combining related people and tags in the user profile improves the results slightly further, leading to a 70:30 ratio between interesting and non-interesting items when explanations are included. In addition, a hybrid people-tag-based recommender has other advantages, such as low proportion of expected items, high diversity of item types, richer explanations, and the simple fact that for some users, recommendations based on people work better, while for others, recommendations based on tags are more effective. Future work should thoroughly examine whether the results presented here can be further improved by means such as integration of other recommenders (e.g., content-based or popularity-based), execution of more sophisticated algorithms (e.g., clustering of people, tags, or items), or optimization of the parameters used by the recommender engine.

7. ACKNOWLEDGMENTS

We thank Sigalit Ur and Tal Daniel for designing and implementing the recommender widget. We are grateful to Shila Ofek-Koifman and Sivan Yogev for many useful discussions.

8. REFERENCES

- [1] Amitay, E., Carmel, D., Har'el, N., Soffer, A., Golbandi, N., Ofek-Koifman, S., & Yogev, S. 2009. Social Search and Discovery using a Unified Approach. *Proc. HYPERTEXT '09*.
- [2] Balabanovic, M. & Shoham, Y. 1997. Fab: Content-based, Collaborative Recommendation. *Commun. ACM* 40, 3 (Mar. 1997), 66-72.
- [3] Bonhard, P. & Sasse, M. A. 2006. Knowing me, Knowing you - Using Profiles and Social Networking to Improve Recommender Systems. *BT Technology Journal* 24, 3 (Jul. 2006), 84-98.
- [4] Burke, R. 2002. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction* 12, 4 (2002), 331-370.
- [5] Carmel, D., Zwerdling, N., Guy I., Ofek-Koifman, S., Har'el N., Ronen, I., Uziel, E., Yogev, S., & Chernov, S. 2009. Personalized Social Search based on the User's Social Network. *Proc. CIKM '09*, 1227-1236.
- [6] Claypool, M., Le, P., Wased, M., & Brown, D. 2001. Implicit Interest Indicators. *Proc. IUI '01*, 33-40.
- [7] Claypool, M., Gokhale, A., Miranda, T., Murnikov, P., Netes, D., & Sartin, M. 1999. Combining Content-Based and Collaborative Filters in an Online Newspaper. *Workshop on Recommender Systems, SIGIR '99*.
- [8] DiMicco, J., Millen, D. R., Geyer, W., Dugan, C., Brownholtz, B., & Muller, M. 2008. Motivations for Social Networking at Work. *Proc. CSCW '08*, 711-720.
- [9] Farrell, S., & Lau T. 2006. Fringe Contacts: People Tagging for the Enterprise. *Workshop on Collaborative Web Tagging, WWW '06*.
- [10] Firan, C. S., Nejd, W., & Paiu, R. 2007. The Benefit of Using Tag-Based Profiles. *Proc. LA-WEB '07*, 32-41.
- [11] Freyne J., Jacovi, M., Guy I., & Geyer W. 2009. Increasing Engagement through Early Recommender Intervention. *Proc RecSys '09*, 85-92.
- [12] Golbeck J. 2006. Generating Predictive Movie Recommendations from Trust in Social Networks. *Proc. 4th Int. Conf. on Trust Management*. Pisa, Italy.
- [13] Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. 1992. Using Collaborative Filtering to Weave an Information Tapestry. *Commun. ACM* 35, 12 (Dec. 1992), 61-70.
- [14] Groh, G., & Ehmig, C. 2007. Recommendations in Taste Related Domains: Collaborative Filtering vs. Social Filtering. *Proc. GROUP '07*, 127-136.
- [15] Guy, I., Jacovi, M., Shahar, E., Meshulam, N., Soroka, V., & Farrell, S. 2008. Harvesting with SONAR: The Value of Aggregating Social Network Information. *Proc. CHI '08*, 1017-1026.
- [16] Guy I., Ronen I., & Wilcox E. 2009. Do You Know? Recommending People to Invite into Your Social Network. *Proc. IUI '09*, 77-86.
- [17] Guy, I., Zwerdling, N., Carmel, D., Ronen, I., Uziel, E., Yogev, S., & Ofek-Koifman, S. 2009. Personalized recommendation of social software items based on social relations. *Proc. RecSys '09*, 53-60.
- [18] IBM Social Software for Business – Lotus Connections: <http://www-01.ibm.com/software/lotus/products/connections/>.
- [19] Kautz, H., Selman, B., & Shah, M. 1997. ReferralWeb: Combining Social Networks and Collaborative Filtering. *Commun. ACM* 40, 3 (Mar. 1997) 63-65.
- [20] Konstas, I., Stathopoulos, V., & Jose, J. M. 2009. On social networks and collaborative recommendation. *Proc. SIGIR '09*, 195-202.
- [21] Lerman, K. 2007. Social Networks and Social Information Filtering on Digg. *Proc. ICWSM '07*.
- [22] Li, X., Guo, L., & Zhao, Y. E. 2008. Tag-based Social Interest Discovery. *Proc. WWW '08*, 675-684.
- [23] Melville, P., Mooney, R. J., & Nagarajan, R. 2002. Content-booster collaborative filtering for improved recommendations. *Proc. AAAI '02*, 187-192.
- [24] Official Digg Blog: <http://blog.digg.com/?p=127>.
- [25] Official YouTube Blog: <http://youtube-global.blogspot.com/2008/06/new-personalized-homepage-and-improved.html>
- [26] Pazzani, M.J., & Billsus D. 2007. Content-based recommendation systems. *The Adaptive Web*, 325-341.
- [27] Ronen, I., Shahar, E., Ur, S., Uziel, E., Yogev, S., Zwerdling, N., Carmel, D., Guy, I., Har'el, N., & Ofek-Koifman, S. 2009. Social networks and discovery in the enterprise (SaND). *Proc. SIGIR '09*, 836.
- [28] Schein, A. I., Popescul, A., Ungar, L. H., & Pennock, D. M. 2002. Methods and Metrics for Cold-start Recommendations. *Proc. SIGIR '02*, 253-260.
- [29] Sen, S., Vig, J., & Riedl, J. 2009. Tagommenders: Connecting Users to Items through Tags. *Proc. WWW '09*, 671-680.
- [30] Sinha, R. & Swearingen, K. 2001. Comparing Recommendations Made by Online Systems and Friends. *DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries*.
- [31] Sinha, R. & Swearingen, K. 2002. The Role of Transparency in Recommender Systems. *CHI Extended Abstracts '02*, 830-831.
- [32] Vatturi, P. K., Geyer, W., Dugan, C., Muller, M., & Brownholtz, B. 2008. Tag-based filtering for personalized bookmark recommendations. *Proc. CIKM '08*, 1395-1396.
- [33] Vig, J., Sen, S., & Riedl, J. 2009. Tagsplanations: Explaining Recommendations using Tags. *Proc. IUI '09*, 47-56.