

A Study of the Benefit of Leveraging Tagging Behavior to Model Users' Interests in del.icio.us

Julia Stoyanovich¹, Sihem Amer Yahia², Cameron Marlow², Cong Yu²

¹Columbia Univ.

New York, NY 10025

jds1@cs.columbia.edu

²Yahoo! Research

New York, NY 10018

{sihem, cameronm, congyu}@yahoo-inc.com

Abstract

Social tagging sites such as Flickr, YouTube and del.icio.us are becoming increasingly popular. Users of these sites annotate and endorse content by tagging, and form social ties with other users by including them into their friendship network. The richness of social context raises the users' expectations with respect to the quality of served content, but also presents a unique opportunity for the design of semantically-enriched recommender systems. This paper presents a variety of methods for producing customized hotlists and evaluates their effectiveness on del.icio.us datasets. We model a user's interest in terms of the tags he uses to annotate content, and in terms of his explicitly stated and derived social ties, and demonstrate how such interest can be leveraged to produce hotlists of very high quality. We also discuss possible research directions and outline strategies for the design of a social tagging recommender system.

Introduction

The increasing popularity of social tagging sites such as Flickr, del.icio.us, and YouTube is justified by the benefit of addressing specific communities of interest when searching for information on the Web. The ability to define explicit social ties with other users in the same social content site raises users' expectations with respect to the quality of served content. For example, in del.icio.us, a social bookmarking and tagging site, users can subscribe to their friends' feeds in order to learn about their latest bookmarked URLs. They can also view hotlists (most popular URLs) as well as browse tags to find related URLs. Similar activities are enabled on photo sites such as Flickr and social networking sites such as Facebook. The unprecedented popularity of these sites is the source of a wealth of user-generated content. The ability to sift through large amounts of content and find the right content to recommend to the right user is a challenging problem that has a big impact on the *survival* of these sites. While Information Retrieval relies on the assumption that content is static and user interests are dynamic and expressed

using keyword search, Information Filtering techniques have been developed to address dynamic content and static user interests (Konstan 2007). We observe that in del.icio.us both content and user interest are dynamic: people tag new content every day. Therefore social tagging sites present a unique opportunity for content recommendation. The ability to model users and their interests in this context is a new challenge.

The simplest form of recommendation is a *hotlist* – a list of most popular items among a set of users in a given period of time. Towards the goal of understanding content recommendation in collaborative tagging sites, we propose to study the generation and personalization of hotlists on del.icio.us datasets. To the best of our knowledge, our work is the first attempt to study content recommendation in social tagging sites.

Our Approach

Hotlists are a mean to expose vital content which has global popularity in the system. While globally popular items usually represent consensus between most users, we experimentally observe that such items only account for a small fraction of any individual user's tagging. We thus look for ways to incorporate individual user's preferences into hotlists computation. We first model the interests of a user as represented by the vocabulary he uses to tag URLs: if a significant portion of the users' tagging actions include the tag *sports*, the user is likely to be interested in sports-related content. This simple observation allows us to replace a single global hotlist by per-tag lists, and to suggest potentially interesting URLs by drawing from one or more per-tag lists in accordance with the user's preferences. We demonstrate that tag-driven customization significantly improves hotlist quality, but its success is still limited by the global aspect.

In the second approach, we propose to model interest using social ties. These ties are either explicitly stated or derived. An example of an explicit social tie is a friendship network. We explore the utility of friendship ties in hotlist generation, and demonstrate experimentally that such ties can indeed be leveraged to generate higher quality hotlists.

Collaborative Filtering (CF) is a popular method

that uses machine learning to determine interest overlap between users based on their behavior such as common ratings of items, or common purchasing and browsing patterns. (Items in collaborative filtering correspond to URLs in our context.) In the present work, we adopt an approach similar to CF and construct a *common interest network* that links two users if the sets of URLs they tagged overlap significantly. We demonstrate how such networks can be used to construct personalized hotlists of very high quality. However, we also observe that using the entire set of URLs tagged by a user as basis for discovering social ties, while leading to high-quality overlap in interest, only applies to a small subset of the users in our user base.

One factor that limits the effectiveness of deriving interest overlap between users in CF is *sparsity*: there are often many more items in the system than any one user is able to rate. This issue is further aggravated in the context of a collaborative tagging site such as del.icio.us, where the set of items corresponds to a potentially infinite set of Internet sites. Sparsity is one of the main reasons why using overlap in items to derive common interest networks is only effective for a subset of del.icio.us users in our experiments. Another important reason is that people rarely agree on everything: you may agree with your mother on cooking, and with your adviser on research, but your adviser’s opinion on food is hardly relevant. We use this idea and demonstrate how tags and item overlap can be combined to construct per-tag common interest networks. Such networks have wider applicability than item-only interest networks, and can be used to construct hotlists of high quality.

We discuss the quality of a hotlist generation method with respect to its scope and coverage. *Scope* refers to the portion of the entire user base to which the method can be applied. It is used to measure the applicability of the hotlist generation method: the wider its applicability – the more users can potentially benefit from it. *Coverage* quantifies the average overlap of the hotlist with the user’s interests. It is a measure of hotlist quality: the higher the coverage – the more representative a hotlist is of the users’ interests. Coverage is related to the concept of *dilution*: the number of users whose opinions are taken into account when generating a hotlist. The more users cast their votes – the less closely the final top-10 list represents the opinion of each of them individually.

This paper makes the following novel contributions:

- We formalize the problem of customized hotlist generation, and propose evaluation metrics to measure the quality of generated hotlists.
- We present a preliminary experimental evaluation of several hotlist generation methods on del.icio.us datasets.
- We discuss the open issues in designing a collaborative tagging recommender system.

Data Model and Problem Statement

We are given a set of users \mathcal{U} , a set of items \mathcal{I} , and a set of tags \mathcal{T} . We will use the following notations:

- $\mathbf{friends}(u)$ is the set of users in \mathcal{U} defined by u to be his friends. This relationship is directional, i.e. the expression below does not always hold:
 $u_1 \in \mathbf{friends}(u_2) \Rightarrow u_2 \in \mathbf{friends}(u_1)$.
- $\mathbf{tags}(u)$ is the set of tags in \mathcal{T} used by user u .
- $\mathbf{taggedBy}(i, u, t)$ is true iff user u has tagged item i with tag t .
 $\mathbf{taggedBy}(i, u)$ is true iff user u has tagged item i with any tag.
- $\mathbf{items}(u, t) = \{i \in \mathcal{I} \mid \mathbf{taggedBy}(i, u, t)\}$ defines all items tagged by u with t .
 $\mathbf{items}(u) = \{i \in \mathcal{I} \mid \mathbf{taggedBy}(i, u)\}$ defines all items tagged by u with any tag.
- $\mathbf{taggers}(i, t) = \{u \in \mathcal{U} \mid \mathbf{taggedBy}(i, u, t)\}$ denotes the set of taggers of an item $i \in \mathcal{I}$ with tag $t \in \mathcal{T}$.
 $\mathbf{taggers}(i) = \{u \in \mathcal{U} \mid \mathbf{taggedBy}(i, u)\}$ denotes the set of taggers of an item $i \in \mathcal{I}$ with any tag.

We use the following terminology to describe a hotlist generation algorithm M :

- *Scope* of M is the set of users $\mathcal{U}_{scope} \subseteq \mathcal{U}$ for whom M generates hotlists.
- *Seed* of M is the set of users $\mathcal{U}_{seed} \subseteq \mathcal{U}$ who are used to generate hotlists for $u \in \mathcal{U}_{scope}$.

The larger the scope of M – the more users can benefit from the hotlist generation method M .

Given a set of items \mathcal{I} , and the seed set \mathcal{U}_{seed} , we define the score of an item $i \in \mathcal{I}$ as the number of users in \mathcal{U}_{seed} who tagged item i :

$$\mathbf{score}(i, \mathcal{U}_{seed}) = |\mathbf{taggers}(i) \cap \mathcal{U}_{seed}| \quad (1)$$

The goal of a method M is to produce a hotlist of items $HList$. In this work, we ignore the ordering of items in the list, and treat the list simply as a subset of \mathcal{I} . $HList$ includes items that have the highest scores from among those tagged by members of \mathcal{U}_{seed} . Without loss of generality, we assume that we generate top-10 hotlists, $|HList| = 10$ for all methods.

We quantify the performance of M in terms of *coverage*: a normalized metric that represents the overlap between the hotlist and the items tagged by the user during the specified time period.

$$\mathbf{coverage}(HList, u) = \frac{|HList \cap \mathbf{items}(u)|}{\min(|\mathbf{items}(u)|, 10)} \quad (2)$$

Coverage of a method M is the average of per-user coverage over all users in \mathcal{U}_{scope} .

$$\mathbf{coverage}(M) = \mathit{AVG}_{u \in \mathcal{U}_{scope}} \mathbf{coverage}(HList, u) \quad (3)$$

We define our problem as producing the hotlist which increases the average coverage.

del.icio.us Data Description

We evaluate the performance of our hotlist generation methods using del.icio.us datasets. del.icio.us is a social tagging site where users contribute content by bookmarking and annotating URLs with tags, and form social ties by declaring their friendship (also referred to as *network*) with other users.

The dataset of tagging actions is very sparse and follows a long tail distribution (Kipp & Campbell ; Golder & Huberman): most URLs are tagged by only a handful of users, and many tags are only used by a few users. Sparse datasets are difficult to process efficiently. We applied the following procedure to reduce the size of the datasets, while still preserving the quality of the generated hotlists. First, we observed that URLs which are rarely tagged stand no chance of contributing to a hotlist. Thus, we removed all URLs that were tagged by fewer than 10 distinct users. Additionally, we removed tagging actions that include uncommon tags: only tags used by at least 4 distinct users are included in our dataset. This cleaning procedure resulted in cutting the tail of URLs and tags. As a result, the dataset was reduced to 27% of its original size.

We evaluate the performance of our hotlist generation methods over a random sample¹ of the cleaned del.icio.us data that includes tagging actions during a consecutive one-month period in 2006. Our cleaned dataset contains 116,177 distinct users who tagged 175,691 distinct URLs using 903 distinct tags, for a total of 2,322,458 tagging actions.

Not all registered del.icio.us users tag URLs. However, for the purposes of our evaluation we focus on taggers – users who contributed at least one tagging action to the cleaned dataset. The terms *users* and *taggers* are used interchangeably in the remainder of this paper.

Evaluating Hotlist Quality

We are now ready to describe the results of our experiments on generating hotlists. All experiments are written as SQL queries over the del.icio.us database which stores users, URLs and tagging actions.

Using Global Popularity

We first consider the quality of hotlists that are based on global popularity of a URL (this is what is referred to as a “hotlist” by most systems). For this method, $\mathcal{U}_{seed} = \mathcal{U}_{scope} = \mathcal{U}$, and the score of a URL i is simply the number of users who tagged that URL: $score(i) = |taggers(i)|$. The top-10 best URLs computed using this method constitute the hotlist, and this hotlist is global, i.e. it is not customized per-user. We refer to this method as **global**.

The average coverage of **global** over all 116,177 users in our experiments is 2.7%. This amount, while quite small, indicates that there is some correlation between the users’ tagging behavior and globally popular URLs.

¹We are not able to reveal our subsetting strategy due to the need to preserve sensitive Yahoo! statistics.

It also argues for improving the method of producing hotlists by accounting for users interests.

Combining Global Popularity and Tags

We now examine a hotlist customization method that defines the interests of a user in terms of tags used by that user. We define the *interest* of a user u for tag t as the fraction of the user’s tagging actions that include t .

$$interest(u, t) = \frac{|items(u, t)|}{|items(u)|} \quad (4)$$

We compute separate top-10 URL lists for each tag $t \in \mathcal{T}$, and experiment with two different ways of using these lists for user hotlist generation. In the first approach, which we call **best_tag**, we identify, for each user u , a single tag from among $tags(u)$ for which $interest(u, t)$ has the highest value, and use the global top-10 URLs for that tag as the user’s hotlist. Ties are broken arbitrarily.

More formally, for a given $t \in \mathcal{T}$,

$$\begin{aligned} \mathcal{U}_{scope} &= \{u \in \mathcal{U} \mid \forall t' \in tags(u), \\ &\quad interest(u, t) \geq interest(u, t')\} \\ \mathcal{U}_{seed} &= \{u \in \mathcal{U} \mid t \in tags(u)\} \end{aligned}$$

The method **best_tag** can be used for all taggers $u \in \mathcal{U}$, and we report the average coverage over all users (116,177) for this experiment. Focusing on the single best tag improves coverage to 8.6% over **global**. A deeper study of the data reveals that using **best_tag** to produce hotlists is most effective for users with comparatively higher values of $interest(u, t)$ for their best tag.

We draw two conclusions from our observation: first, that accounting for tagging behavior when computing hotlists does improve hotlist quality, and second, that a more general method for hotlist selection is needed, particularly for cases where no tag can be identified that clearly dominates a user’s interest.

We thus propose another method, called **dominant_tags**, where we identify taggers who have a strong interest in 1 or several tags, and then combine the best URLs from the per-tag top-10 lists into a single top-10 custom hotlist. For a given tag $t \in \mathcal{T}$,

$$\begin{aligned} \mathcal{U}_{scope} &= \{u \in \mathcal{U} \mid t \in tags(u) \wedge interest(u, t) > thresh\} \\ \mathcal{U}_{seed} &= \{u \in \mathcal{U} \mid t \in tags(u)\} \end{aligned}$$

In the current set of experiments, we say that a user u has a strong interest in a tag t if $interest(u, t) > 0.3$. This threshold was determined empirically, and needs further validation in a future study. With an interest threshold set to 0.3, a user can have at most 3 dominant tags. If more than one tag passes the interest threshold, we draw an equal number of items from the top-10 list that corresponds to each dominant tag. So, if two tags t_1 and t_2 both pass the interest threshold for user $u \in$

# dominant tags	$ \mathcal{U}_{scope} $	coverage
1	36,736	10.4%
2	16,452	14.4%
3	6,466	17.8%

Table 1: Effect of the number of dominant tags on the performance of `dominant_tags`

interest	$ \mathcal{U}_{scope} $	coverage
30%	36,736	10%
40%	31,391	11%
50%	25,703	13%
60%	23,927	13%
70%	20,943	14%
80%	19,704	14%
90%	19,392	14%
100%	19,347	14%

Table 2: Effect of interest on the performance of `dominant_tags`

\mathcal{U}_{scope} , the final list *HLIST* consists of the top-5 entries from *HLIST*₁ and the top-5 entries from *HLIST*₂. For users with 3 dominant tags, we choose the top-3 URLs from each *HLIST*₁, *HLIST*₂, and *HLIST*₃ and build a top-9 hotlist.

Table 1 lists the partition of the users in our dataset by the number of tags for which they have strong interest, and presents performance of `dominant_tags` for these users. Note that `dominant_tags` has less than perfect scope: the total number of users partitioned this way is smaller than the entire user base; 49% of del.icio.us users in our dataset have no dominant tags.

The first row of Table 1 reports coverage as an average over the 36,736 users who have one dominant tag. For those users, the generated hotlist constitutes 10% of all URLs tagged by each user, on average. This number increases when users have 2 and 3 dominant tags. Clearly, the more dominant tags there are for a user – the better the coverage. However, the higher the number of dominant tags, the smaller the number of users, i.e. the more limited the scope of applicability of the hotlist generation method.

We now turn our attention to the relationship between strength of a user’s interest for a tag, and the coverage of the `dominant_tags` method. We do this for users with a single dominant tag (row 1 in Table 1). Table 2 summarizes our findings: as expected, the stronger the user’s affinity for a tag, the better the coverage, but also the more limited the scope.

We argue here that Table 2 motivates a richer hotlist generation strategy: while coverage increases with the increasing affinity for a tag, a plateau is reached as interest reaches 60%. Hotlists generated by `dominant_tags` take into account a user’s interests, as derived from his linguistic choices. However, items in the tag-specific top-10 lists still represent consensus of

a very large group of users: $|\mathcal{U}_{seed}|$ is as large as 29,712 for the most popular tag in our experiments, i.e. 29,712 distinct users in our sample used this tag. For users with two or three dominant tags, multiple top-10 lists are combined, further increasing the size of \mathcal{U}_{seed} .

In large seed sets, opinions and interests of individual users are approximated. The larger the seed – the coarser the approximation. We call this effect *dilution*. To minimize dilution, given a user $u \in \mathcal{U}_{scope}$, a hotlist generation strategy needs to identify the seed set \mathcal{U}_{seed} that is both representative of the interests of the user u , and focused on that user’s interests.

In the remainder of this section, we seek to reduce dilution by considering explicit and implicit social ties between users during hotlist generation.

Computing Hotlists using Friendship

The goal of this experiment is to explore the utility of the explicitly defined friendship network in del.icio.us for hotlist generation. Of 116,177 taggers in our experiments, 36,248 (31%) also participate in the friendship network in our dataset, and we term such users *friendly taggers*.

Choosing the friendship network as \mathcal{U}_{seed} is justified by the fact that del.icio.us users tend to pay attention to their friends tagging actions, which influences their own. In particular, users can subscribe to their friends’ feeds and get notified whenever one of their friends tags a new URL.

For each friendly tagger, we generate a customized hotlist by drawing 10 URLs with highest popularity from among URLs tagged by his friends. We refer to this hotlist generation method as `friends`. More formally, we define the scope of `friends` as the set of friendly taggers:

$$\mathcal{U}_{scope} = \{u \in \mathcal{U} | \exists f \in \mathcal{U} \wedge f \in \text{friends}(u)\}$$

We next define the seed for a fixed user $u \in \mathcal{U}_{scope}$ as follows:

$$\mathcal{U}_{seed} = \{f \in \text{friends}(u)\}$$

We focus on a *random subset* of friendly taggers, 4,644 in all, for our experiments in this section, and find the coverage of `friends` to be 42.9%, a significant improvement over `global`, which was 3.3% for the sample of friendly taggers in our experiments. Note that the $avg(|\mathcal{U}_{seed}|) = 4$ for our sample, multiple orders of magnitude less than in the previous sections!

We now explore how tagging can be leveraged to deduce interest overlap among users, and how such overlap can be used to generate customized hotlists of very high quality. We report two experiments: in the first, a tie is derived between two users if the sets of URLs they tag overlap significantly; in the second, we enrich the set of derived ties by considering tag-specific overlap in URLs.

agreement	$ \mathcal{U}_{scope} $	$avg(\mathcal{U}_{seed})$	coverage
30%	1382	169.4	61.47%
50%	913	136.9	72.82%

Table 3: Effect of the agreement threshold on the effectiveness of `url_interest`

Interest as Overlap in URLs

In this experiment we compute a *URL Interest Network* by considering the overlap in tagged URLs between users. We quantify *agreement* between users u_1 and u_2 as the fraction of URLs tagged by u_1 that were also tagged by u_2 .

$$\text{agr}(u_1, u_2) = \frac{|\text{items}(u_1) \cap \text{items}(u_2)|}{|\text{items}(u_1)|} \quad (5)$$

Note that agreement is directional. If $\text{agr}(u_1, u_2)$ is above a certain *agreement threshold*, we will use URLs tagged by u_2 to derive the hotlist for u_1 . We refer to this method as `url_interest`. More formally, we first define the scope of the hotlist generation method as the set of users whose agreement with at least one other user in \mathcal{U} is above the agreement threshold:

$$\mathcal{U}_{scope} = \{u \in \mathcal{U} | \exists f \in \mathcal{U} \wedge \text{agr}(u, f) > \text{thresh}\}$$

We next define the seed for a fixed user $u \in \mathcal{U}_{scope}$ as the set of users with whom u agrees:

$$\mathcal{U}_{seed} = \{f \in \mathcal{U} \wedge \text{agr}(u, f) > \text{thresh}\}$$

Table 3 summarizes the effectiveness of `url_interest` in terms of scope and coverage. We observe that while the method achieves very good coverage, it is very limited in its scope: only 1382 users can benefit from customized hotlists if 30% agreement is required. The scope is further limited to 913 users for a minimum agreement of 50%. Note also that $avg(|\mathcal{U}_{seed}|)$ is lower for the 50% threshold. We believe this to be a case of lower dilution (fewer users in the \mathcal{U}_{seed}) leading to higher coverage. However, further experimental validation is needed to better understand the effect of dilution on scope and coverage of `url_interest`.

The limited scope of `url_interest` is due to the fact that strong agreement of the kind required by this method is rather uncommon. In order to include an edge between two users in the interest network, we require that they agree on at least 30% of the tagged URLs over-all. We observe that, while agreement of this kind over all interests may be rare, people more commonly agree with others on only a part of their interests. We explore this idea in the final experimental section.

Interest as Overlap in URLs and in Tags

In this section, we propose to combine interest in a tag, as explored in `dominant_tags`, with agreement based

method	$ \mathcal{U}_{scope} $	$avg(\mathcal{U}_{seed})$	coverage
<code>dominant_tags</code>	1235	26,855.5	17.2%
<code>tag_url_interest</code>	1235	227.4	81.7%
<code>url_interest</code>	205	202.9	84.5%

Table 4: Relative performance of `dominant_tags`, `tag_url_interest`, and `url_interest`

on URL overlap, as in the previous section, to construct a *tag-URL-interest network*. We use tag interest to generate a global partitioning of tagging actions. We then search for URL-agreement within these partitions. We call this method `tag_url_interest`, and propose to use it for the taggers who show strong interest in one or more tags (see Equation 4). Thus, we define the scope of `tag_url_interest` as the set of taggers with strong interest in a tag, and with URL agreement, for that tag, with at least one other user. We first define a tag-specific version of agreement as:

$$\text{agr}(u_1, u_2, t) = \frac{|\text{items}(u_1, t) \cap \text{items}(u_2, t)|}{|\text{items}(u_1, t)|} \quad (6)$$

The scope of `tag_url_interest` is then, for a fixed tag $t \in \mathcal{T}$:

$$\mathcal{U}_{scope} = \{u \in \mathcal{U} | t \in \text{tags}(u) \wedge \text{interest}(u, t) > \text{int_thresh} \wedge \exists f \in \mathcal{U} \wedge \text{agr}(u, f) > \text{agr_thresh}\}$$

We next define the seed for a fixed $u \in \mathcal{U}_{scope}$, and for a fixed tag $t \in \mathcal{T}$ as the set of users who are in tag-url-agreement with u and who have used t in their tagging:

$$\mathcal{U}_{seed} = \{f \in \mathcal{U} | t \in \text{tags}(f) \wedge f \in \text{agr}(u, f) > \text{agr_thresh}\}$$

We evaluated the effectiveness of `tag_url_interest` on a subset of users in our experiments: we choose users with strong interest in exactly 2 tags. Out of 16,452 users with strong interest in 2 tags, 1235 were in the scope of `tag_url_interest`. We note, that while the scope of the current method is still limited, it greatly exceeds the scope of `url_interest`: only 205 users in the scope of `tag_url_interest` where also in the scope of `url_interest`.

Table 4 summarizes the relative performance of `dominant_tags`, `url_interest` and `tag_url_interest` for the users in \mathcal{U}_{scope} of `tag_url_interest`. The first line summarizes performance of `dominant_tags` for users in `tag_url_interest`. Because \mathcal{U}_{scope} of `tag_url_interest` is a proper subset of \mathcal{U}_{scope} of `dominant_tags`, we report performance of `dominant_tags` for all users in `tag_url_interest`. In this table, we use 30% as the threshold for both interest and URL agreement, wherever appropriate.

We note that `tag_url_interest` significantly outperforms `dominant_tags` in terms of coverage. However, for users who are in both `tag_url_interest` and

`url_interest`, the latter does better. This represents 205 users for whom `tag_url_resource` achieves an 81.7% coverage while `url_resource` achieves an 84.5% coverage. This reinforces the idea that accounting for agreement over all URLs is stronger than agreement on one tag at a time. This result also supports our dilution hypothesis: $avg(|U_{seed}|)$ is smallest for `url_interest`, followed by `tag_url_interest`.

Summary and Discussion

In the present work we demonstrate that tagging behavior of users in a social tagging site can be leveraged to derive implicit social ties between users. We also show that such ties serve as a good indicator of a user's interest. Clearly, hotlists generated by observing social ties are of higher relevance than those produced by global hotlist generation strategies. Ultimately, we would like to devise a general strategy for the design of *social tagging recommender systems*. We classify the main challenges towards this goal along two axes: *interaction with users* and *core recommendation strategy*, and discuss each of them in turn in the remainder of this section.

Interaction with Users

Explanations An important concern in recommender systems is the ability to explain why a particular item is recommended to a user. It has been shown in the past that providing an explanation for recommendations improves the trust of the user in the system (Konstan 2007). While numerical values (average item ratings) have been used in the past by recommenders such as Amazon and Launch, they only provide a coarse explanation of recommendations. Some systems support their recommendations by statements about previous performance such as "MovieLens has predicted correctly 80% of the time for you". Complex explanations such as two dimensional statistics, e.g., correlation and variance, often do not work. Instead of stating "People who have tagged this URL, also tagged these others", we propose to make explanations more precise by providing semantic information of the form "These URLs are being recommended to you by people who share the same interest as you in Sports" or "This URL is being recommended to you by your friends". The ability to offer such explanations is facilitated by the use of social ties, some of which may be labeled with the tags that were used to derive them.

Sparsity and Incentives Machine learning methods (e.g., Collaborative Filtering (CF)) have been used extensively in state of the art recommendation systems. One factor that makes generating recommendation particularly difficult is sparsity: the menu of items from which to recommend is large, while user judgments about these items are few. To address sparsity, machine learning systems often go through one or multiple approximation steps, which in turn makes explaining recommendations very difficult. In our work, we also

observe that tagging data is sparse and follows a long tail distribution: many URLs are only tagged once or a handful of times, and many tags are only used by one or very few distinct users. However, our approach towards dealing with sparsity is different from that used in CF: we aim to leverage explicit and implicit social ties when making and explaining recommendation. A method like ours also has the potential to solicit and incorporate user feedback into the recommendation process. Explanations and the use of feedback to improve user experience serve as incentives to the user to contribute to the system.

Core Recommendation Strategy

Users Our experiments demonstrate that there is significant heterogeneity in users' tagging behavior, and thus in the types of social ties that can be derived for different users. Applicability and effectiveness of a recommendation strategy to a particular user is directly influenced by these social ties. We argue for partitioning the users into categories that fit their behavior and interests for purposes of recommendation. Moreover, a user may benefit from receiving recommendations based on different strategies.

Interest: External Factors There are different reasons why a user would bookmark a URL and use specific tags to label it. Our evaluation of hotlist recommendation strategies shows that a user's interest in a URL is influenced by his social ties (explicit friendship), many or all of which may be external: formed in real life, in another social content site, etc. Another external influence over a user's interest in a URL or in a tag is the general popularity of that URL or tag. Bursts in interest have been observed and studied in the past (Cohen & McCallum). Indeed, some tags appear at certain periods of time and indicate a trend in the general public or among a community of users. For example, during holiday time, people tend to visit travel sites more often than at other times. Wars, health, and news also tend to follow similar trends, while other tags, such as cooking or professional photography, have a steadier popularity. From the point of view of recommendation, this calls for a better understanding of tags and justifies the need to classify them in a way that accounts for their popularity. We envision a system that involves the user in the choice and classification of relevant tags in a time-sensitive manner.

Persistence of Interest and Agreement Collaborative Filtering is based on a strong underlying assumption that people who agreed in the past tend to agree again in the future, and conversely that items that are similar will continue to be similarly liked or disliked by a given user. Consequently, existing CF systems do little to detect or predict changes in user preferences. During our experimental evaluation we found that the assumption of persistence of interest and agreement does not hold in del.icio.us: derived social ties, particularly

those based on resource agreement, were rarely conserved from one time period to the next. We identified the short lifespan of URLs as the main reason: most URLs are only tagged actively for a period of 1-2 weeks. This phenomenon makes it essential to consider and incorporate time into any successful social tagging recommendation strategy. Our current efforts focus on expanding URL lifetime by grouping multiple URLs into sets, and drawing recommendations from these sets.

Related Work

The idea of motivating participation by displaying the value of contribution is characteristic of collaborative reviewing sites (Rashid *et al.* 2006) but has received little attention in the study of collaborative tagging sites. Impact of reviews has been studied extensively in the e-commerce arena, and it has been shown that reviews impact sales (Chen, Dhanasobhon, & Smith 2007; Chevalier & Mayzlin 2006; Ghose & Ipeirotis 2006). Most of the current scientific literature dealing with user-contributed reviews concerns text analysis to distill reviews (Popescu & Etzioni 2005), sentiment detection (Pang & Lee 2004; 2005), and the impact of reviews on product sales (Bickart & Schindler 2001; Ghose & Ipeirotis 2007; Kim *et al.* 2006; Liu *et al.* 2007). In contrast, very little has been done to extract value in collaborative tagging systems and provide participation incentives to users.

Collaborative Filtering is a popular method which is based on using machine learning to derive and compare user profiles in order to determine overlap of interest between users. A user profile is built from explicit or implicit data. The system can explicitly ask users to rate an item on a numerical scale, or to rank a collection of items from most to least favorite. The system can also record items that a user browsed or purchased, and analyze item viewing times. CF compares data collected from the user to similar data collected from others and calculates a list of recommended items for the user. Several methods have been developed to address data sparsity. Most of them (item-based and user-based) rely on statistical approximation (Bell, Koren, & Volinsky 2007; Park & Pennock 2007).

Our method differs from CF in that we express a user's interest qualitatively rather than quantitatively, using tags and derived ties. In (Agichtein, Brill, & Dumais 2006), it is shown that ranking in Web search can be improved by incorporating user behavior. Similarly, we show that incorporating tagging behavior improves ranking in producing hotlists. This motivates the need to better understand the principles behind designing a recommender tagging system, as was briefly discussed in (Golder & Huberman). According to a study of del.icio.us tagging practices described in (Kipp & Campbell), tagging exhibits self-organizing patterns. In our work we explore how users can be classified into groups based on their tagging behavior, and how such

groups can be used to improve the quality of recommended hotlists.

References

- Agichtein, E.; Brill, E.; and Dumais, S. 2006. Improving web search ranking by incorporating user behavior information. In *SIGIR06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, 19–26. New York, NY, USA: ACM Press.
- Bell, R.; Koren, Y.; and Volinsky, C. 2007. Modeling relationships at multiple scales to improve accuracy of large recommender systems. In *KDD '07: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, 95–104. New York, NY, USA: ACM Press.
- Bickart, B., and Schindler, R. M. 2001. Internet forums as influential sources of consumer information. *Journal of Interactive Marketing* 15(3):31–40.
- Chen, P.; Dhanasobhon, S.; and Smith, M. 2007. All reviews are not created equal: The disaggregate impact of reviews and reviewers at amazon.com. In *Proceedings of the 28th Annual International Conference on Information Systems (ICIS2007)*. To appear. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=918083.
- Chevalier, J. A., and Mayzlin, D. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* 345–354.
- Cohen, W., and McCallum, A. Information extraction from the world wide web. In *KDD03: The Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Ghose, A., and Ipeirotis, P. 2006. Designing ranking systems for consumer reviews : The impact of review subjectivity on product sales and review quality. In *Proceedings of the Workshop on Information Technology and Systems (WITS)*.
- Ghose, A., and Ipeirotis, P. G. 2007. Designing ranking systems for consumer reviews : The economic impact of customer sentiment in electronic markets. In *Proceedings of the International Conference on Decision Support Systems (ICDSS)*.
- Golder, S. A., and Huberman, B. A. The structure of collaborative tagging systems. *Information Dynamics Lab, HP Labs*.
- Kim, S.-M.; Pantel, P.; Chklovski, T.; and Pennacchiotti, M. 2006. Automatically assessing review helpfulness. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP-06)*, 423–430.
- Kipp, M. E., and Campbell, D. G. Patterns and inconsistencies in collaborative tagging systems: An examination of tagging practices. *Faculty of Information and Media Studies University of Western Ontario*. Available from <http://dlist.sir.arizona.edu/1704/01/KippCampbellASIST.pdf>.

Konstan, J. A. 2007. Introduction to recommender systems. In *SIGIR07: Proceedings of the 30th Annual International ACM SIGIR Conference*.

Liu, J.; Cao, Y.; Lin, C. Y.; Huang, Y.; and Zhou, M. 2007. Low-quality product review detection in opinion summarization. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, 334–342.

Pang, B., and Lee, L. 2004. A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL '04: Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, 271. Morristown, NJ, USA: Association for Computational Linguistics.

Pang, B., and Lee, L. 2005. Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL '05: Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, 115–124. Morristown, NJ, USA: Association for Computational Linguistics.

Park, S.-T., and Pennock, D. M. 2007. Applying collaborative filtering techniques to movie search for better ranking and browsing. In *KDD '07: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, 550–559. New York, NY, USA: ACM Press.

Popescu, A.-M., and Etzioni, O. 2005. Extracting product features and opinions from reviews. In *HLT '05: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, 339–346. Morristown, NJ, USA: Association for Computational Linguistics.

Rashid, A. M.; Ling, K.; Tassone, R. D.; Resnick, P.; Kraut, R.; and Riedl, J. 2006. Motivating participation by displaying the value of contribution. In *CHI '06: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 955–958. New York, NY, USA: ACM Press.