Improving FolkRank With Item-Based Collaborative Filtering

Jonathan Gemmell, Thomas Schimoler, Maryam Ramezani,
Laura Christiansen, Bamshad Mobasher
Center for Web Intelligence
School of Computing, DePaul University
Chicago, Illinois, USA
{jgemmell, tschimo1, mramezani, lchris10, mobasher}@cdm.depaul.edu

ABSTRACT

Collaborative tagging applications allow users to annotate online resources. The result is a complex tapestry of interrelated users, resources and tags often called a folksonomy. Folksonomies present an attractive target for data mining applications such as tag recommenders. A challenge of tag recommendation remains the adaptation of traditional recommendation techniques originally designed to work with two dimensional data. To date the most successful recommenders have been graph based approaches which explicitly connects all three components of the folksonomy.

In this paper we speculate that graph based tag recommendation can be improved by coupling it with item-based collaborative filtering. We motive this hypothesis with a discussion of informational channels in folksonomies and provide a theoretical explanation of the additive potential for item-based collaborative filtering. We then provided experimental results on hybrid tag recommenders built from graph models and other techniques based on popularity, user-based collaborative filtering and item-based collaborative filtering.

We demonstrate that a hybrid recommender built from a graph based model and item-based collaborative filtering outperforms its constituent recommenders. Furthermore the inability of the other recommenders to improve upon the graph-based approach suggests that they offer information already included in the graph based model. These results confirm our conjecture. We provide extensive evaluation of the hybrids using data collected from three real world collaborative tagging applications.

1. INTRODUCTION

Collaborative tagging has emerged as a popular method for organizing and sharing online content with user-defined keywords. Delicious[1], Flickr[2] and Last.fm[3] are among the most popular destinations on the Web allowing users to annotate bookmarks, digital photographs and music respectively. Other less popular tagging applications serve niche communities enabling users to tag blogs, business documents or scholarly articles.

[1]delicious.com
[2]www.flickr.com
[3]www.last.fm

RecSys'09, October 22–25, 2008, New York City, New York.

At the heart of collaborative tagging is the post; a user describes a resource with a set of tags. A collection of posts results in a complex network of interrelated users, resources and tags commonly referred to as a folksonomy [16]. Users are able to navigate this network free from a rigid conceptual hierarchy.

Despite the freedom users enjoy, the size of a folksonomy often hampers the userŠs exploration. Data mining applications such as recommenders can assist the user by reducing a burdensome number of items to a smaller collection related the user's interests. In this work we focus on tag recommendation, the suggestion of tags during the annotation process.

Tag recommendation reduces the cognitive effort from generation to recognition. Users are therefore encouraged to tag more frequently, apply more tags to a resource, reuse common tags and use tags the user had not previously considered. User error is reduced by eliminating capitalization inconsistencies, punctuation errors, misspellings and other discrepancies. The final result is a cleaner denser dataset that is useful in its own right or for further data mining applications.

Despite the richness offered by folksonomies, they also present unique challenges for tag recommenders. Traditional recommendation strategies, often developed to work with two dimensional data, must be adapted to work with the three dimensional nature of folksonomies. Otherwise they risk disregarding potentially useful information. To date the most successful tag recommenders are graph-based models, which exploits the user-defined links between the users, resources and tags.

In this work we propose augmenting the graph based approach with item-based collaborative filtering. We offer a discussion of information channels in folksonomies to motivate this proposal. The graph based model covers the user-resource, user-tag, and resource-tag channels. Item-based collaborative filtering, on the other hand, focuses on tags previously applied by the user to resources similar to the query resource. It therefore includes resource-resource information not explicitly contained in the graph model. Additionally, the user-tag information utilized by item-based collaborative filtering is more oriented to query resource.

We construct hybrid tag recommenders composed of the graph models and other techniques including popularity models, user-based collaborative filtering and item-based collaborative filtering. The graph based recommender coupled with item-based collaborative filtering produces better results than either produce alone, strengthening our theory that that item-based collaborative filtering contains information that is absent in the graph based model. Moreover the other hybrids do not improve upon the graph based model suggesting that the information they contain are already adequately represented by the graph based approach.

The rest of this paper is organized as follows. In Section 2 we describe related works. A brief survey of the tag recommenders we employ in our experiments is given in Section 3. The use of hybrid recommenders is motivated in Section 4 where we discuss informational channels in folksonomies. Section 5 details how tag recommenders may be compounded to produce hybrid recommenders. Our experimental evaluation is presented in Section 6, including a description of our datasets, our methodology and a discussion of our findings. Finally in Section 7 we present our conclusions and lay a foundation for future work.

2. BACKGROUND AND RELATED WORK

The term *folksonomy* was coined by [28], a play on *folk* and *taxonomy*. While the term is new, [29] argues that collaborative tagging in merely a renaissance of manual indexing. However, the scope and connectivity of the Internet permits tagging to rise to a level heretofore unrealized.

In [16] the attractiveness of tagging is outlined: serendipitous browsing, a low entry cost, utilizing the wisdom of the crowd, and a sense of community. Moreover, he argues that tagging allows objects to be categorized under multiple tags, unfettered from traditional taxonomies. He also discusses two obstacles: tag ambiguity in which a tag has several meanings and tag redundancy in which several tags have the same meaning.

As collaborative tagging applications have gained in popularity researchers have explored and characterized the tagging phenomenon. In [15] and [10] the authors studied the information dynamics of Delicious, one of the most popular folksonomies. The authors discussed how tags have been used by individual users over time and how tags for an individual resource stabilize over time. In [15] the authors provide an overview of the phenomenon and offer reasons why both folksonomies and taxonomies will have a place in the future of information access.

There have been many recent research investigations into recommendation within folksonomies. Unlike traditional recommender systems which have a two-dimensional relation between users and items, tagging systems have a three dimensional relation between users, tags and resources. Recommender systems can be used to recommend each of the dimensions based on one or two of the other dimensions. In [26] the authors apply user-based and item-based collaborative filtering to recommend resources in a tagging system and uses tags as an extension to the user-item matrices. Tags are used as context information to recommend resources in [19] and [18].

In [13] user-based collaborative filtering is compared to a graph-based recommender based on the PageRank algorithm for tag recommendation. The authors in [11] use association rules to recommend tags and introduce an entropy-based metric to define how predictable a tag is. In [14] the title of a resource, the posts of a resource and the user's vocabulary are used to recommend tags.

User-defined tags and co-occurrence are employed by [24] to recommend tags to users on Flickr. The assumption is that the user has already assigned a set of tags to a photo and the recommender uses those tags to recommend more tags. The authors in [6] have completed a similar study and introduce a classification for tag recommendation. Probabilistic models have been used in recommendation in folksonomies in [20] and [30]. Moreover, [20] uses Probabilistic Latent Semantic Analysis for resource discovery and [30] uses single aspect PLSA for tag recommendation.

Previously, in [8, 9], we demonstrated how tag clusters serving as coherent topics can aid in the personalization of search and navigation. Further support for the utility of clustering is offered in [4] where improvement in search through clustering is theorized. In [7] we adapted K-Nearest Neighbor for tag recommendation and showed incorporating user tagging habits into recommendation can improve K-Nearest Neighbor.

General criteria for a good tagging system including high coverage of multiple channels, high popularity and least-effort are presented in [31]. They categorize tags as content-based tags, context-based tags, attribute tags, subjective tags, and organizational tags and use a probabilistic method to recommend tags. In [2] the authors propose a classification algorithm for tag recommendation. Semantic tag recommendation systems in the context of a semantic desktop are explored in [1]. Clustering to make real-time tag recommendation is developed in [25].

3. TAG RECOMMENDATION

Here we first provide a model of folksonomies, then review several common recommendation techniques which we employ in our evaluation. A folksonomy can be described as a four-tuple:

$$D = \langle U, R, T, A \rangle \tag{1}$$

where, U is a set of users; R is a set of resources; T is a set of tags; and A is a set of annotations, represented as user-tag-resource triples:

$$A \subseteq \{\langle u, r, t \rangle : u \in U, r \in R, t \in T\} \tag{2}$$

A folksonomy can, therefore, be viewed as a tripartite hyper-graph [17] with users, tags, and resources represented as nodes and the annotations represented as hyper-edges connecting a user, a tag and a resource.

Aggregate projections of the data can be constructed, reducing the dimensionality but sacrificing information [22]. The relation between resources and tags, RT, can be formulated such that each entry, $RT(r, t)$, is the weight associated with the resource, r, and the tag, t. This weight may be binary, merely showing that one or more users have applied that tag to the resource. In this work we assume $RT(r, t)$ to be the number of users that have applied t to the r:

$$RT_{tf}(r, t) = |\{a = \langle u, r, t \rangle \in A : u \in U\}| \tag{3}$$

Analogous two-dimensional projections can be constructed for *UT* in which the weights correspond to users and tags, and *UR* in which the weights correspond to users and resources.

Many authors have attempted to exploit the data model for recommendation in folksonomies. In traditional recommendation algorithms the input is often a user, u, and the output is a set of items, I. Tag recommendation differs in that the input is both a user and a resource. The output remains a set of items, in this case a set of recommended tags, T_r. Given a user-resource pair, the recommendation set is constructed by calculating a weight for each tag, $w(u, r, t)$, and recommending the top n tags.

3.1 Popularity Based Approaches

We consider two popularity based models which rely on the frequency a tag is used. **PopRes** ignores the user and relies on the popularity of a tag within the context of a particular resource. We define the resource based popularity measure as:

$$w(u, r, t) = \frac{|\{a = \langle u, r, t \rangle \in A : u \in U\}|}{|\{a = \langle u, r, t \rangle \in A : u \in U, t \in T\}|} \tag{4}$$

PopUser, on the other hand, ignores the resource and focuses on the frequency of a tag within the user profile. We define the user based popularity measure as:

$$w(u, r, t) = \frac{|\{a = \langle u, r, t\rangle \in A : r \in R\}|}{|\{a = \langle u, r, t\rangle \in A : r \in R, t \in T\}|} \quad (5)$$

Popularity based recommenders require little online computation. Models are built offline and can be incrementally updated. However both these models focus on a single channel of the folksonomy and may not incorporate otherwise relevant information into the recommendation.

3.2 User-Based Collaborative Filtering

User-based K-nearest neighbor is a commonly used recommendation algorithm in Information Retrieval that can be modified for use in folksonomies. Applications may model users by recency, authority, linkage or vector space models. In this work we focus on the vector space model [21] and describe the user as a vector over either the tag space or the resource space.

KNN_UT models the user, u, as a vector over the set of tags where the weight in each dimension corresponds to the occurrence of the tag in the user profile as it is defined by the two dimensional projection $UT(u, t)$. Other methods may be used to model the user, such as a vector over the set of resources or a combination of tags and resources. Several techniques may be used to calculate the similarity between vectors such as Jaccard similarity or cosine similarity [27]. In this work we rely on cosine similarity.

Using the similarity measure a neighborhood, N, of the k most similar users is constructed such that they have all previously annotated the query resource, r. A weight for each tag is calculated as:

$$w(u, r, t) = \frac{\sum_n^N sim(u, n) * d(n, r, t)}{k} \quad (6)$$

where $d(n, r, t)$ is 1 if the neighbor, n, has annotated the query resource, r, with the tag t. Otherwise it is 0.

Traditional user-based collaborative filtering requires a comparison between the query user and every other user. However, since the adapted algorithm considers only those users that have annotated the query resource, the number of similarities to calculate is drastically reduced. The popularity of resources in folksonomies follows the power law and the great majority of resources will benefit from this reduced reduction in computation, while a few will require additional computational effort. As a result the algorithm scales well with large datasets.

However, since the algorithm relies on the collaboration of other users it may be the case that a tag cannot be recommended because it does not appear in a neighbor's profile. While the personalization offered by user-based filtering is an important component for the recommender, it lacks the ability to reflect the habits and patterns of the larger crowd.

3.3 Item-Based Collaborative Filtering

KNN_RT models resources as a vector over the tag space. Give a resource and a tag, we define the weight as the entry of the two dimensional projection, $RT(r, t)$, the number of times r has been tagged with t. When a user selects a resource to annotate, the cosine similarity between it and every resource in the user profile is calculated. A neighborhood of the k most similar resources, S, is then constructed. We then define the item-based collaborative filtering measure as:

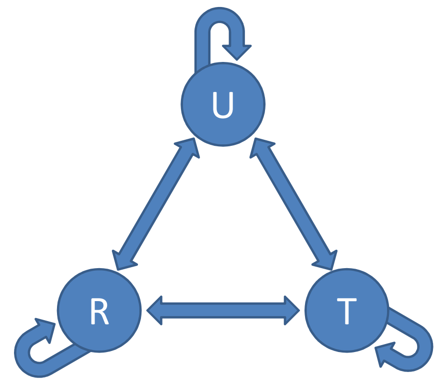

Figure 1: Informational channels of a folksonomy.

$$w(u, r, t) = \frac{\sum_s^S sim(s, r) * d(u, s, t)}{k} \quad (7)$$

where $d(u, s, t)$ will equal 1 if the user has applied t to s and 0 otherwise. This recommender focuses entirely on the user's tagging habits. Unlike the user-based filtering methods, it may be able to identify tags that are common to the user but rarely used by others. However, it lacks the ability to discover relevant tags from other users. Depending on the size of the user profile, this recommender will also scale well to larger datasets, particularly if the resource-resource similarity matrix if calculated offline.

3.4 FolkRank

FolkRank was proposed in [12]. It computes a PageRank vector from the tripartite graph of the folksonomy. This graph is generated by regarding $U \cup R \cup T$ as the set of vertices. Edges are defined by the three two-dimensional projections of the hyper-graph, RT, UR and UT.

If we regard the adjacency matrix of this graph, W, (normalized to be column-stochastic), a damping factor, d, and a preference vector, p, then we iteratively compute the PageRank vector, w, in the usual manner: $w = dAw + (1 - d)p$.

However due to the symmetry inherent in the graph, this basic PageRank may focus too heavily on the most popular elements. The *FolkRank* vector is taken as a difference between two computations of PageRank: one with and one without a preference vector. Tag recommendations are generated by biasing the preference vector towards the query user and resource [13]. These elements are given a substantial weight while all other elements have uniformly small weights.

PageRank has proven to be one of the top performing tag recommenders. However, it imposes steep computational costs.

4. INFORMATIONAL CHANNELS OF FOLKSONOMIES

The model of a folksonomy suggests several informational channels which may be exploited by data mining applications such as tag recommenders. The relation between users, resources and tags generate a complex network of interrelated items as shown in Figure 1.

The channel between resources and tags reveals a highly descriptive model of the resources. The accumulation of many users' opinions (often numbered in the thousands or millions) results in a rich-

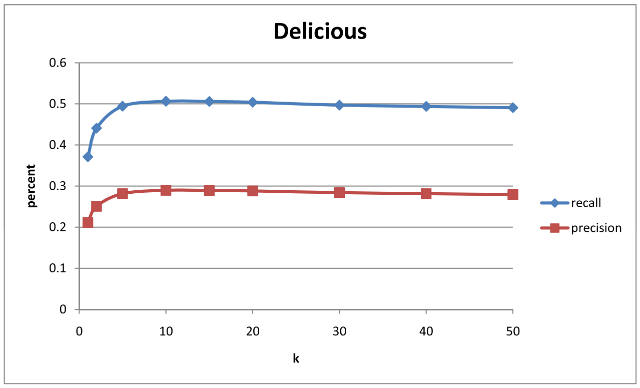

Figure 2: The effect of k in *KNN_UT* on recall and precision for a recommendation set of 5 tags. Users are modeled as a vector over the tag space.

ness which taxonomies are unable to approximate. Conversely the tags themselves are characterized by the resources to which they have been assigned.

As users annotate resource with tags they define their interests in as much as they describe a resource. The user-tag channel therefore reveals the users' interests and provides opportunities for data mining algorithms to offer a high degree of personalization. Likewise a user may be defined by the resources which he has annotated as in the user-resource channel.

These primary channels can be used to produce secondary informational channels. The user-user channel can be constructed by modeling users as a vector of tags or as a vector of resources and applying a similarity measure such as cosine similarity. Many variations exist. However the result reveals a network of users that can be explored directly or incorporated into further data mining approaches. The resource-resource and tag-tag channels provide similar utility, presenting navigational opportunities for users to explore similar resources or neighborhoods of tags.

The success of tag recommenders hinge on their ability to incorporate all of these informational channels. A simple recommender such as *PopRes* focuses only on the tag-resource channel, whereas *PopUser* includes only the information between tags and users.

Collaborative filtering techniques include additional channels but increase the computational overhead. *KNN_UT* discovers a set of neighbors, thereby covering the user-user channel. It then focuses on tags those neighbors applied to the query resource covering the user-resource and resource-tag channels. *FolkRank*, on the other hand, explicitly defines the relation between users, resources and tags in its adjacency matrix. While *FolkRank* has proven to be among most effective tag recommenders, augmenting it with algorithms that incorporate complimentary informational channels may improve its performance.

5. HYBRID RECOMMENDERS

The multiple informational channels of folksonomies present an attractive target for hybrid recommenders. Hybrids combine several recommenders together to produce a new recommender. The constituent recommenders are freed from the burden of the covering all the available informational channels and may instead focus on only a few. The hybrid then ties these recommenders together. A successful hybrid creates a synergistic blend of its constituent parts producing superior results that they could not achieve alone.

In this paper we focus on weighted hybrid recommenders [5] which combine pairs of recommenders in a linear model. Each model is trained separately. Given a user, u, and a resource, r, the hybrid queries both components for each tag in the folksonomy. The results is $W(u, r, t)$ which contains the weights for all tags. In order to ensure that weight assignments for each recommendation approach are on the same scale, we normalize the weights in $W(u, r, t)$ to 1 producing $W'(u, r, t)$.

Originally, these weights were used to select the top n items for the recommendation set. In this case, however, the weights are combined in a linear model as:

$$w(u, r, t) = \beta * w'_a(u, r, t) + \alpha * w'_b(u, r, t) \quad (8)$$

where $\beta = 1 - \alpha$. These coefficients are used to control the contribution of the two recommenders. When α is set to 0, recommender a acts alone. In the case that α is set to 0.5, each recommender contributes equally to the final weight. For each hybrid, α must be empirically tuned to achieve the maximum synergy between the components. The tags are then resorted by the new weight, and the top n tags are recommended for the annotation.

6. EXPERIMENTAL EVALUATION

In this section we describe the methods used to gather and preprocess our datasets. Our testing methodology is outlined. We provide a discussion of how we tuned variables for each algorithm and describe the experiments on the weighted hybrid recommenders. Finally, we discuss our observations.

6.1 Datasets

Folksonomy	Delicious (5%)	Citeulike	Bibsonomy
Users	7,665	2,051	357
Resources	15,612	5,376	1,738
Tags	5,746	3,343	1,573
Posts	720,788	42,278	19,909
Annotations	2,762,235	105,873	54,848

Table 1: Datasets

We provide an extensive evaluation of the hybrid recommenders using data from three real collaborative tagging applications: Delicious, Citeulike, and Bibsonomy.

6.1.1 *P-Core Processing*

By P-core processing users, resources and tags are removed from the dataset in order to produce a residual dataset that guarantees each user, resource and tag occur in at least p posts [3]. Here we define a post to include a user, a resource, and every tag the user has applied to the resource.

By removing infrequent users, resources and tags noise in the data is reduced. Uncommon items whether they be tags used by only a few users, unpopular resources, or inactive users are eliminated from consideration. Because of their scarcity these are the very items likely to confound recommenders. Moreover by eliminating infrequent items the size of the dataset is dramatically reduced allowing the application of data mining techniques that might otherwise be computationally impractical.

6.1.2 *Delicious*

Delicious is a popular collaborative tagging application in which users annotate URLs. On 10/19/2008, 198 of the most popular tags were taken from the user interface. For each of these tags the 2,000

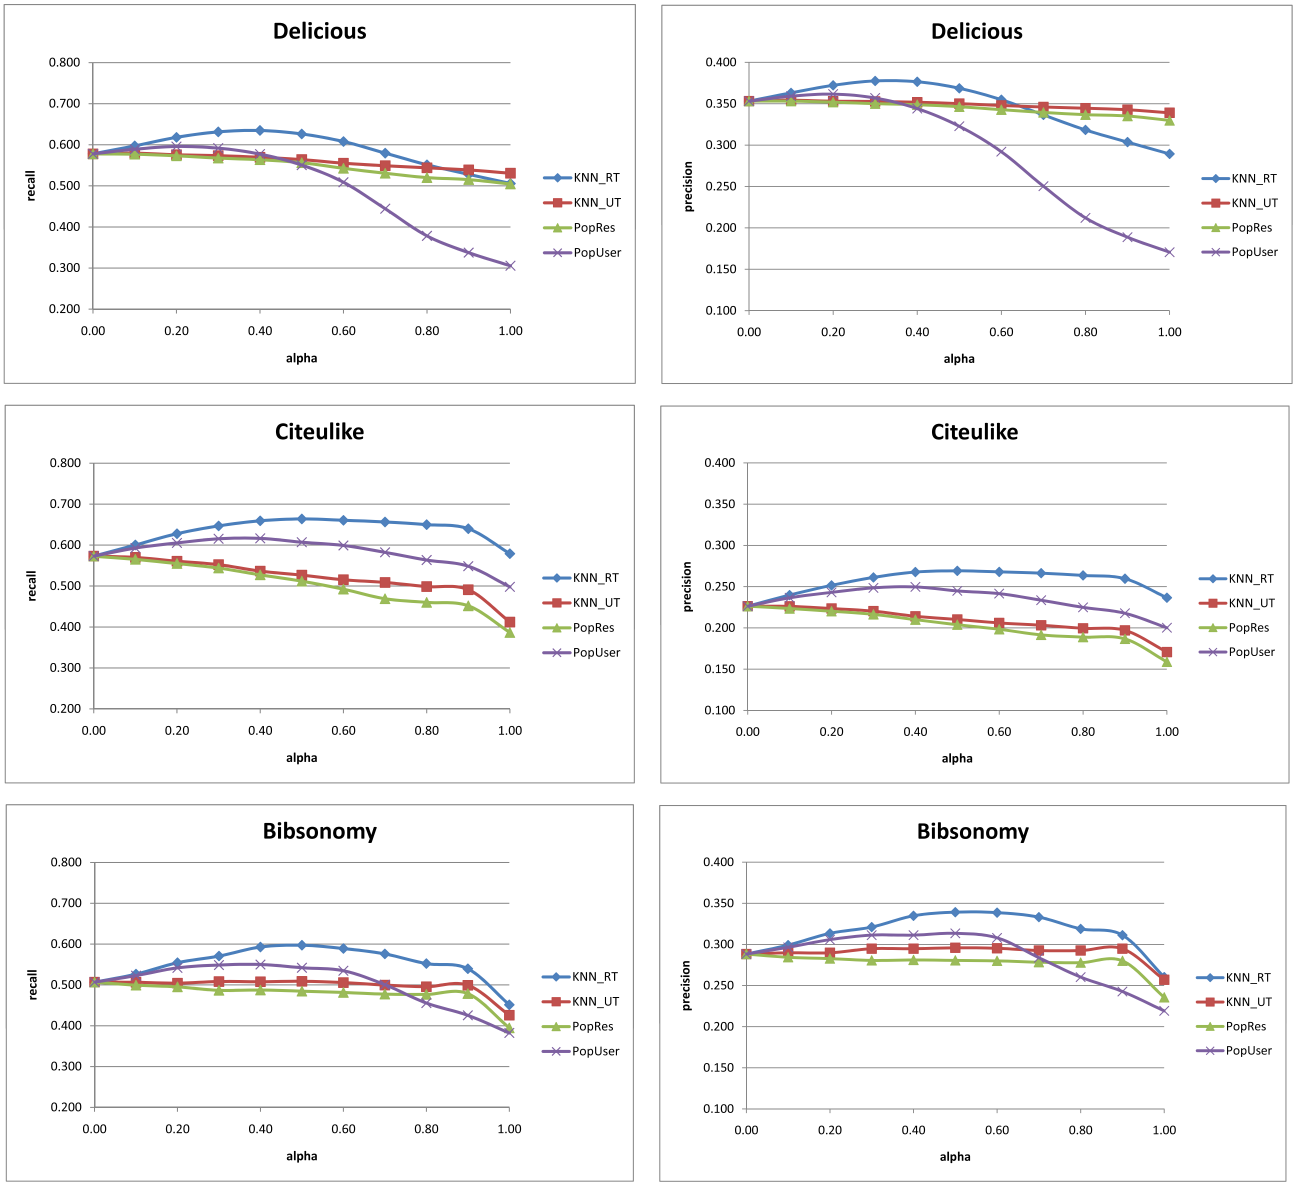

Figure 3: The effect of *alpha* on the hybrid recommenders on the Delicious, Citeulike and Bibsonomy datasets. Results are shown using recall and precision on a recommendation set of five tags.

most recent annotations including the contributors of the annotations were collected. The social network for these contributors was explored recursively collecting 524,790 usernames.

From 10/20/2008 to 12/15/2008 the complete profiles of the users were collected. Each user profile consisted of a collection of annotations including the resource, tags and date of the original bookmark. The top 100 most prolific users were visually inspected; twelve were removed from the data because their annotation count was many orders of magnitude larger than other users and were therefore suspected to be Web-bots.

Due to memory and time constraints, 5% of the user profiles was randomly selected. Still this dataset remains far larger than either the following Bibsonomy or Citeulike datasets. Experiments on larger samplings reveal near identical trends for several of the tag recommendation strategies. Some tag recommendation techniques such as *FolkRank* are so computational intensive that larger samplings of the data are not feasible. In order to best compare the recommenders, the 5% sampling was used on all reported experiments. A P-core of 20 was taken from the sample and is reported in Table 1.

6.1.3 Citeulike

Citeulike is a popular online tool used by researchers to manage and discover scholarly references. They make their dataset freely available to download[4]. On 2/17/2009 the most recent snapshot was downloaded. The data contains anonymous user ids and posts for each user including resources, the date and time of the posting

[4] www.citeulike.org/faq

and the tags applied to the resource. A P-core of 5 was taken. The characteristics of the dataset are described in Table 1.

6.1.4 *Bibsonomy*

This dataset was provided by Bibsonomy[5] for use in the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD) 2009 Challenge. Bibsonomy was originally launched as a collaborative tagging application allowing users to organize and share scholarly references. It has since expanded its scope allowing users to annotate URLs.

The data includes all public bookmarks and publication posts of Bibsonomy until 2009-01-01. The data was cleaned by removing all characters which are neither numbers nor letters from tags. Additionally the system tags *imported*, *public*, *systemimported*, *nn* and *systemunfiled* where removed. A P-core of 5 was used. Table 1 relates the features of the dataset.

6.2 Experimental Methodology

We have adopted the test methodology as described in [13]. In this approach, called *LeavePostOut*, a single post is randomly removed from each user's profile. The training set is then comprised of the remaining posts, while the test set contains one post per user. Each test case consists of a user, u, a resource, r, and all the tags the user has applied to that resource. These tags, T_h, are analogous to the holdout set commonly used in Information Retrieval. The tag recommendation algorithms accept the user-resource pair and return an ordered set of recommended tags, T_r.

For evaluation we adopt the common recall are precision measures as is common in Information Retrieval. Recall measures the percentage of items in the holdout set that appear in the recommendation set. It is a measure of completeness and is defined as:

$$r = |T_h \cap T_r|/|T_h| \quad (9)$$

Precision measures the percentage of items in the recommendation set that appear in the holdout set. It measures the exactness of the recommendation algorithm and is defined as:

$$p = |T_h \cap T_r|/|T_r| \quad (10)$$

For each evaluation metric the average value is calculated across all test cases.

6.3 Experimental Results

Here we present our experimental results beginning with the tuning of variables. The experiments with user-based collaborative filtering require the tuning of k, the number of neighbors.

Figure 2 shows the relation between k and the evaluation metrics recall and precision for a recommendation set of size 5. The Delicious dataset was used for this experiment. As k increases so does recall and precision. However this improvement suffers from diminishing returns until a k of 50 offers little more benefit than a k of 20. This trend was observed for K-Nearest Neighbor experiments in the other two datasets as well. As such, all *KNN_UT* experiments were completed using a k of 20.

Item-based collaborative filtering also requires the tuning of k, in this case the number of similar resources in the user profile to include in the neighborhood. After empirical analysis we found 15 to produce the best performance on all datasets.

Figure 3 shows the tuning of α for the hybrid recommenders. Each hybrid is a linear combination of *FolkRank* and one of the

[5]www.bibsonomy.org

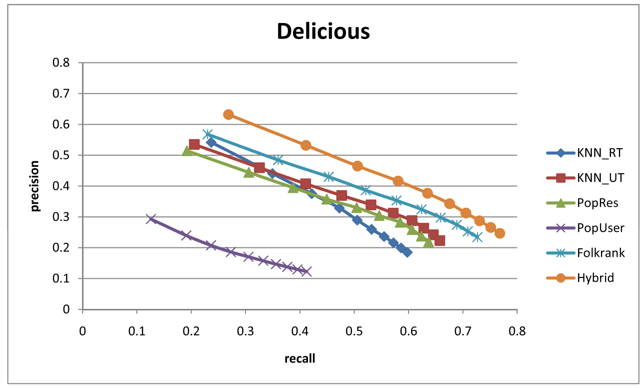

Figure 4: A comparison of tag recommender techniques in Delicious.

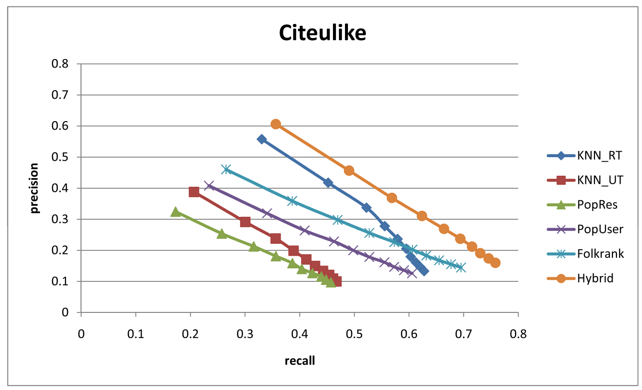

Figure 5: A comparison of tag recommender techniques in Citeulike.

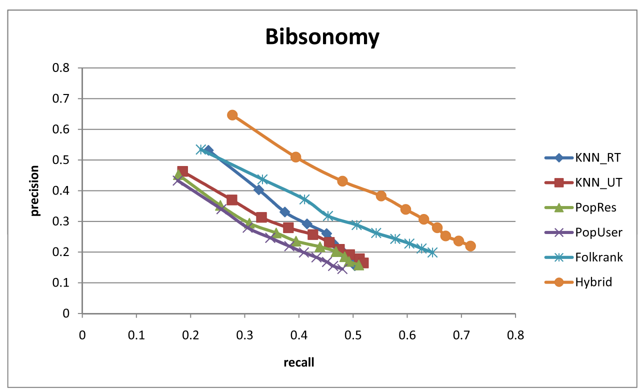

Figure 6: A comparison of tag recommender techniques in Bibsonomy.

other four recommenders. The left hand side of each graph shows the hybrid recommenders when α is set to 0 in which case *FolkRank* dominates the hybrid. As α increases more weight is given to the other recommenders until finally when α reaches 1, *FolkRank* plays no part in the recommendation.

For all datasets, item-based collaborative filtering contributes to recall and precision of its hybrid. For example in the Delicious experiment when α is set to 0.4, recall for a recommendation set of five tags is 6% higher than *FolkRank* achieves alone and 13% higher than *KNN_RT* achieves alone.

In the Delicious experiments, a hybrid built with *PopUser* offers a slight improvement, while it has a more dramatic improvement on Citeulike. These observations reveal that the personalization of the user-tag channel strongly incorporated into *KNN_RT* and *PopUser* offers information lacking in *FolkRank*. While *PopUser* boosts all of the user's tags, *KNN_RT* focuses on tags related to the resource being annotated accounted for its increased performance. On the other hand *PopRes* does not appear to provide any additional benefit to *FolkRank*. Indeed, *FolkRank* contains this information in the utilization of the RT matrix.

These two results reveal that the weights given to the query resource and query user in the *FolkRank* algorithm achieve different results. The weight applied to the resource immediately activates tags strongly associated with the resource. The result is similar to that achieved in *PopRes*, hence *PopRes* offers little assistance to its hybrid. However, the weight applied to the query user disperses through the graph activating all of the user's tags relevant or irrelevant to user's present context. *KNN_RT*, on the other hand, focuses on tags applied to resources similar to the query. Hence, it includes the resource-resource channel missing in *FolkRank*. The hybrid is able to be personalized but also be more context specific.

KNN_UT does not appear to offer any additional information that *FolkRank* did not already contain, even though it includes user-resource information in the neighborhood selection, user-resource information in the cosine similarity and resource-tag information in the recommendation step. This reveals that the way in which the informational channels is equally important. Additionally *KNN_UT* selects neighbors that are similar to the query user, utilizing the user-user channel. However, this channel does not appear to be beneficial to tag recommendation.

After analysis of the effect of α on the hybrids we selected the best α for the *FolkRank-KNN_RT* hybrid. For Delicious we used an α of 0.4. For Citeulike and Bibsonomy used an α of 0.5. Figures 4 through 6 compare tag recommenders along with the hybrid. Recall and precision are plotted for recommendation sets of size one through ten. For all datasets the hybrid outperforms its constituent parts.

We also observe a difference in the effect that constituent recommenders have across the datasets. Delicious users tag Web pages and their topics cover a wide array of topics. Citeulike users tag scholarly articles and often focus on their area of expertise. In fact we can see in Figures 4 and 5 the dramatic difference between *PopRes* and *PopUser*.

In Delicious *PopRes* outperforms *PopUser*, whereas in Citeulike the opposite is true. The user's focus on a narrow subject area in Citeulike make the user-tag channel a informative predictor, whereas the topic variety in the profiles of Delicious users make the resource-tag channel more reliable.

This analysis is underscored by the success *KNN_RT* hybrid has on the Delicious datasets where *PopUser* hybrid fairs poorly. Because *KNN_RT* focuses on those tags applied to resources similar to the query resource it offers context appropriate tags. In Citeulike, where users have a narrow focus, this context provides little additional benefit and the *PopUser* hybrid performs nearly as well as the *KNN_RT* hybrid. Bibsonomy users tags both citations and web pages; its results fall between those of the other two datasets.

7. CONCLUSIONS

We have demonstrated that tag recommenders may be combined to form weighted hybrids that perform better than either performs alone. Moreover *FolkRank* one of the most successful tag recommenders to date can be augmented with item-based collaborative filtering to produce superior results. The resource-resource and personalized user-resource channels covered by item-based collaborative filtering compliment the channels utilized by *FolkRank*. The inability of other recommenders to improve upon *FolkRank* provides evidence that *FolkRank* sufficiently incorporates the informational channels covered by those recommenders.

Future work will involve investigating alternative hybrid tag recommenders. New recommenders that cover other informational channels will be considered. Finally, alternative methods for hybridizing recommenders will be explored.

8. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation Cyber Trust program under Grant IIS-0430303 and a grant from the Department of Education, Graduate Assistance in the Area of National Need, P200A070536.

9. REFERENCES

[1] B. Adrian, L. Sauermann, and T. Roth-Berghofer. Contag: A semantic tag recommendation system. In T. Pellegrini and S. Schaffert, editors, *Proceedings of I-Semantics' 07*, pages pp. 297–304. JUCS, 2007.

[2] P. Basile, D. Gendarmi, F. Lanubile, and G. Semeraro. Recommending smart tags in a social bookmarking system. In *Bridging the Gep between Semantic Web and Web 2.0 (SemNet 2007)*, pages 22–29, 2007.

[3] V. Batagelj and M. Zaveršnik. Generalized cores. *Arxiv preprint cs/0202039*, 2002.

[4] G. Begelman, P. Keller, and F. Smadja. Automated Tag Clustering: Improving search and exploration in the tag space. *Proceedings of the Collaborative Web Tagging Workshop at WWW*, Volume 6, 2006.

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

[6] N. Garg and I. Weber. Personalized, interactive tag recommendation for flickr. In *RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems*, pages 67–74, New York, NY, USA, 2008. ACM.

[7] J. Gemmell, T. Schimoler, M. Ramezani, and B. Mobasher. Adapting k-nearest neighbor for tag recommendation in folksonomies. *Intelligent Techniques for Web Personalization & Recommender Systems*, 2009.

[8] J. Gemmell, A. Shepitsen, B. Mobasher, and R. Burke. Personalization in Folksonomies Based on Tag Clustering. *Intelligent Techniques for Web Personalization & Recommender Systems*, 2008.

[9] J. Gemmell, A. Shepitsen, B. Mobasher, and R. Burke. Personalizing navigation in folksonomies using hierarchical tag clustering. In *Proceedings of the 10th international conference on Data Warehousing and Knowledge Discovery*. Springer, 2008.

[10] S. A. Golder and B. A. Huberman. Usage patterns of collaborative tagging systems. *Journal of Information Science*, 32(2):198, 2006.

[11] P. Heymann, D. Ramage, and H. Garcia-Molina. Social tag prediction. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 531–538, New York, NY, USA, 2008. ACM.

[12] A. Hotho, R. Jaschke, C. Schmitz, and G. Stumme. Information retrieval in folksonomies: Search and ranking. *Lecture Notes in Computer Science*, 4011:411, 2006.

[13] R. Jaschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag Recommendations in Folksonomies. *LECTURE NOTES IN COMPUTER SCIENCE*, 4702:506, 2007.

[14] M. Lipczak. Tag recommendation for folksonomies oriented towards individual users. In *Proceedings of the ECML/PKDD 2008 Discovery Challenge Workshop, part of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2008.

[15] G. Macgregor and E. McCulloch. Collaborative tagging as a knowledge organisation and resource discovery tool. *Library Review*, 55(5):291–300, 2006.

[16] A. Mathes. Folksonomies-Cooperative Classification and Communication Through Shared Metadata. *Computer Mediated Communication, (Doctoral Seminar), Graduate School of Library and Information Science, University of Illinois Urbana-Champaign, December*, 2004.

[17] P. Mika. Ontologies are us: A unified model of social networks and semantics. *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(1):5–15, 2007.

[18] R. Y. Nakamoto, S. Nakajima, J. Miyazaki, S. Uemura, and H. Kato. Investigation of the effectiveness of tag-based contextual collaborative filtering in website recommendation. In *Advances in Communication Systems and Electrical Engineering*, pages 309–318. Springerlink, 2008.

[19] R. Y. Nakamoto, S. Nakajima, J. Miyazaki, S. Uemura, H. Kato, and Y. Inagaki. Reasonable tag-based collaborative filtering for social tagging systems. In *WICOW '08: Proceeding of the 2nd ACM workshop on Information credibility on the web*, pages 11–18, New York, NY, USA, 2008. ACM.

[20] A. Plangprasopchok and K. Lerman. Exploiting social annotation for automatic resource discovery. *CoRR*, abs/0704.1675, 2007.

[21] G. Salton, A. Wong, and C. Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.

[22] C. Schmitz, A. Hotho, R. Jaschke, and G. Stumme. Mining association rules in folksonomies. In *Proc. IFCS 2006 Conference*, pages 261–270. Springer, 2006.

[23] B. Sigurbjörnsson and R. van Zwol. Flickr tag recommendation based on collective knowledge. pages 327–336, 2008.

[24] B. Sigurbjörnsson and R. van Zwol. Flickr tag recommendation based on collective knowledge. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 327–336, New York, NY, USA, 2008. ACM.

[25] Y. Song, Z. Zhuang, H. Li, Q. Zhao, J. Li, W.-C. Lee, and C. L. Giles. Real-time automatic tag recommendation. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 515–522, New York, NY, USA, 2008. ACM.

[26] K. H. L. Tso-Sutter, L. B. Marinho, and L. Schmidt-Thieme. Tag-aware recommender systems by fusion of collaborative filtering algorithms. In *SAC '08: Proceedings of the 2008 ACM symposium on Applied computing*, pages 1995–1999, New York, NY, USA, 2008. ACM.

[27] C. Van Rijsbergen. *Information Retrieval*. Butterworth-Heinemann Newton, MA, USA, 1979.

[28] T. Vander Wal. Folksonomy definition and wikipedia. *vanderwal. net*, 2005.

[29] J. Voss. Tagging, Folksonomy & Co-Renaissance of Manual Indexing? *Arxiv preprint cs/0701072*, 2007.

[30] R. Wetzker, W. Umbrath, and A. Said. A hybrid approach to item recommendation in folksonomies. In *ESAIR '09: Proceedings of the WSDM '09 Workshop on Exploiting Semantic Annotations in Information Retrieval*, pages 25–29, New York, NY, USA, 2009. ACM.

[31] Z. Xu, Y. Fu, J. Mao, and D. Su. Towards the semantic web: Collaborative tag suggestions. *Collaborative Web Tagging Workshop at WWW2006, Edinburgh, Scotland, May*, 2006.