On the Effect of Group Structures on Ranking Strategies in Folksonomies

Fabian Abel
IVS – Semantic Web Group
Leibniz University Hannover
Appelstr. 4 D-30167 Hannover
abel@kbs.uni-hannover.de

Nicola Henze
IVS – Semantic Web Group
Leibniz University Hannover
Appelstr. 4 D-30167 Hannover
henze@kbs.uni-hannover.de

Daniel Krause
IVS – Semantic Web Group
Leibniz University Hannover
Appelstr. 4 D-30167 Hannover
krause@kbs.uni-hannover.de

Matthias Kriesell
Department of Mathematics
University of Hamburg
Bundesstraße 55 D-20146 Hamburg
kriesell@math.uni-hamburg.de

ABSTRACT

Folksonomies have shown interesting potential for improving information discovery and exploration. Recent folksonomy systems investigate on the use of tag assignments, which combine Web resources with annotations (tags), and the users that have created the annotations. In this paper, we investigate on the effect of grouping resources, i.e. creating sets of resources, and using this additional structure for search in folksonomies. Our experiments show that the quality of search result ranking can be significantly improved by introducing and exploiting the grouping of resources (one-tailed t-Test, level of significance $\alpha = 0.05$).

Categories and Subject Descriptors

H.3.3 [**Information Systems**]: Information Search and Retrieval; H.4.m [**Information Systems**]: Miscellaneous; H.5.3 [**Information Systems**]: Information interfaces and presentationGroup and Organization Interfaces

General Terms

Social Search

Keywords

Social Media, Search, Ranking, Folksonomies, Web 2.0

1. INTRODUCTION

The success of social systems in the Web is quite obvious, for example with popular systems like Flickr[1], YouTube[2], Blogger[3] and others more. Systems like these allow users to share photos, broadcast own videos, or blog about topics they are interested in. They document impressively that the active contribution of users to the creation process of content, and the possibility to share content immediately with (fellow) users, is highly requested by Web users. The *tagging* activity is one of the important characteristics of these systems: With tagging, a user adds freely chosen words that come into his or her mind when watching or using some content. The result of a single tagging activity is a binding between a user, a resource, and the respective keywords that this user assumes relevant for the resource. The evolving set of such bindings *user-resource-tag* is called a folksonomy [12]. By nature, the folksonomy is highly dynamic, and an important characteristic is that the tags assigned by users are not bound to any controlled vocabulary but contribute to a growing set of words.

The so-far developed folksonomy systems all have in common that the set of user-resource-tag bindings is hardly structured any further. Del.icio.us[4] allows to structure tags by grouping them into so-called *bundles*, and bibsonomy [6] allows to structure users by formation of user groups. So far, nobody has investigated the effect of structuring the *resource dimension* in folksonomies, and, to the best of our knowledge, no present ranking algorithm takes further structure within the respective sets of users, tags or resources into account in order to improve the performance.

We have realized an appealing Web 2.0 application that enables users to easily construct groups of Web content that they consider interesting for some topic. GroupMe! users can group arbitrary Web resources like videos, news feeds, images, etc. Within a GroupMe! group these resources are visualized according to their media type – e.g. videos can directly be played within a group, news feeds list their latest items, etc. – so that content of groups is easy to grasp. GroupMe!'s tagging functionality allows users to annotate both, resources and groups. Hence, whenever resources are annotated, this is done in context of a group.

The immediate benefit of the GroupMe! approach is that

[1]http://flickr.com
[2]http://youtube.com
[3]http://blogger.com

WWW2008, April 21–25, 2008, Beijing, China.
.

[4]http://del.icio.us

we are now able to see Web resources in a context, namely the group context: Web resources which were previously not related at all now have in common that they belong to some group which defines a common context. Together with tagging, we can even further specify this relation between the members of a group: The group's tags are likely to be relevant for the members of the group, and vice versa. This is particularly interesting for the discovery of Web resources: the resource's context gives us means to find relevant tags even if the resource itself has not been tagged at all. We were able to show that the GroupMe! approach improves the recall of retrieval (see [2]).

Groups of content provide us with a database of *hand-picked resources* for certain topics, which are specified by the group and its tags. Presumably, these resources are of high relevance for the topic – in comparison to search results lists – as a subject is screening the search results and decides which to add to the group, and which not. In this paper, we investigate how to make use of this database of hand-picked resources, and how to exploit the grouping structure on resources in order to improve the quality of ranking strategies in folksonomies. We benchmark our investigation against a popular ranking strategy in folksonomies, the FolkRank algorithm [7]. It turns out that the grouping structure significantly improves the quality of ranking.

The paper is organized as follows: In the next section, we describe the GroupMe! system in more detail. Section 3 defines the folksonomy model, and explains the necessary extensions of this model to reflect the grouping mechanism. Ranking algorithms, for both the folksonomy model with and without groups, are discussed in Section 4. The analysis of our experimental findings are presented thereafter. The paper ends with a comparison to related work, and the conclusion.

2. GROUPME! TAGGING SYSTEM

The GroupMe! system[5] [1, 2], online since July 2007, is a classic tagging system similar to del.icio.us or BibSonomy as it allows users to search for resources or tag their own resources as well as the resources of other users.However, GroupMe! has a novel function that, as far as we know, no other present tagging system offers: Organizing resources by arranging them in groups. Users can use groups to structure their resources according to different topics, or to get a better overview of multiple resources as every resource is visualized in a multimedia-based fashion: Pictures are displayed as thumbnails, videos can be played immediately "within the group", and RSS feeds are previewed by displaying the recent headlines. Hence, users can get a quick overview of multiple resources by using the GroupMe! system. An example is given in Figure 1 where a group *Travel to the WWW 2008 conference* is displayed. All relevant information, like a link to the official website, an image of some conference hotel, a short video tutorial for Chinese language, and other content can be seen at a glance.

GroupMe! handles groups just like other resources, i.e. groups can be tagged like ordinary resources and groups can be included in other groups, allowing to build a hierarchical group structure. E.g. in Figure 1 the WWW conference group also contains another group called *Business trip to Beijing*, which informs about common useful issues regarding business trips to Beijing.

Creating groups is very easy as users can put new resources into a group via simple *drag & drop* operations. GroupMe! includes interfaces for some major sources of resources. Currently, we provide search interfaces for Google, Flickr, and GroupMe! resources. Compared to the task of tagging, which requires a user to think of an appropriate keyword and entering this, grouping of resources does not require any keyword interaction and can be performed very fast. This enables us to gather more information from a higher percentage of the users than in traditional tagging systems. This additional information can be used in different ways, e.g. for extracting sub- and superclass relationships between tags, or for inheriting tags through the group hierarchy. The second method will be used in this paper to improve the keyword-based search in folksonomies.

3. FOLKSONOMIES

The term *folksonomy*, introduced by Thomas Vander Wal in 2004 [12], defines a taxonomy, which evolves over time when users (the *folks*) annotate resources with freely chosen keywords. Folksonomies can be divided into *broad* folksonomies, which allow different users to assign the same tag to the same resource, and *narrow* folksonomies, in which the same tag can be assigned to a resource only once [11]. Formal models of a folksonomy are e.g. presented in [4, 13]. They are based on bindings between users, tags, and resources. According to [6] a folksonomy is defined as follows:

Definition 1. A *folksonomy* is a quadruple $\mathbb{F} := (U, T, R, Y)$, where:

- U, T, R are finite sets of instances of *users*, *tags*, and *resources*, respectively, and
- Y defines a relation, the *tag assignment*, between these sets, that is, $Y \subseteq U \times T \times R$.

In [15], tag assignments are furthermore attributed with a timestamp and Hotho et al. also embed relations between tags into the formal folksonomy model [6]. In order to simplify the formalization we do not include these features. GroupMe! introduces groups as a new dimension in folksonomies.

Definition 2. A *group* is a finite set of resources.

A group is a resource as well. Groups can be tagged or arranged in groups, which effect hierarchies among resources. In general, tagging of resources within the GroupMe! system is done in context of a group. Figure 2 presents a basic GroupMe! tagging scenario, in which users u_1 and u_2 have grouped resources r_{1-3} into g_1 and g_2, and have tagged both, resources and groups with keywords t_{1-3}. The tag assignment (u_1, t_2, r_2, g_1) in Figure 2 describes that user u_1 has annotated resource r_2 in context of group g_1 with tag t_2. If users assign tags to a group, which is itself not contained in a group, then the group context information is not available ($\rightarrow (u_2, t_2, g_2, \varepsilon)$). Hence, a GroupMe! folksonomy is formally characterized via Definition 3 (cf. [1]).

Definition 3. A GroupMe! folksonomy is a 5-tuple $\mathbb{F} := (U, T, \breve{R}, G, \breve{Y})$, where:

- U, T, R, G are finite sets that contain instances of users, tags, resources, and groups, respectively,

[5]http://groupme.org

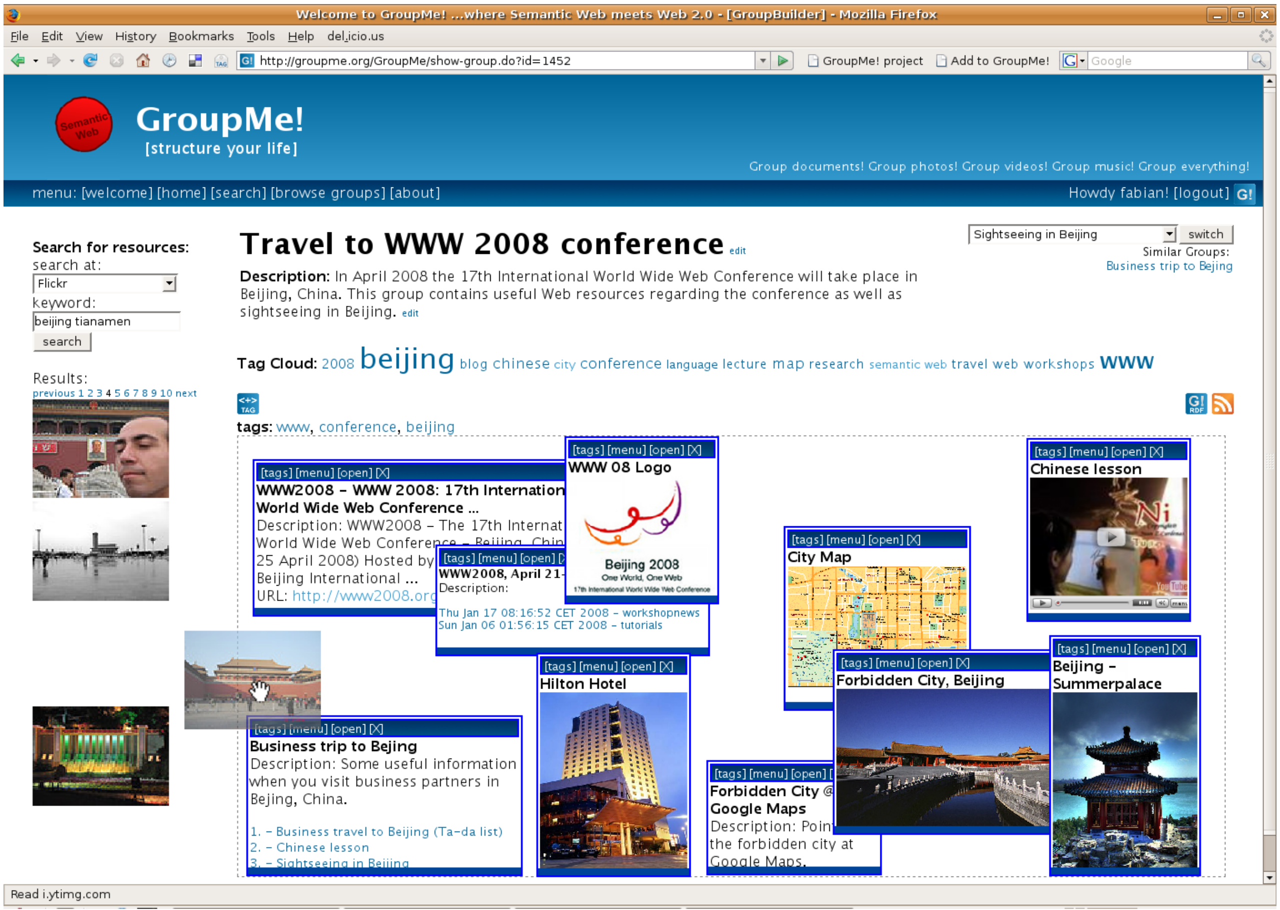

Figure 1: Screenshot of the GroupMe! system: Constructing groups via *drag & drop*

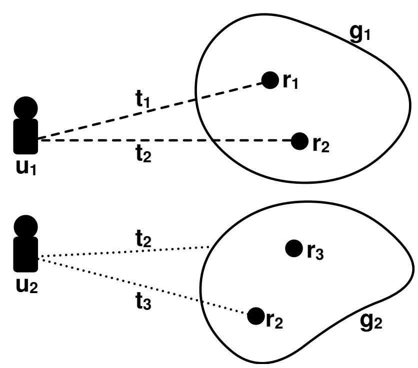

Figure 2: Scenario in which two users assign tags to resources in context of different groups.

- $\breve{R} = R \cup G$ is the union of the set of resources and the set of groups, and
- $\breve{Y}$ defines a *GroupMe! tag assignment*: $\breve{Y} \subseteq U \times T \times \breve{R} \times (G \cup \{\varepsilon\})$, where ε is a reserved symbol for the *empty group context*, i.e. a group that is not contained in another group when it gets tagged by a user.

In comparison to traditional folksonomies (see Definition 1), in which relations between tags mainly rely on their co-occurrences (i.e. two tags are assigned to the same resource), a GroupMe! folksonomy gains new relations between tags:

1. A relation between tags assigned from (possibly) different users to different resources, where the resources are contained in the same group.
2. A relation between tags assigned to a group g and tags assigned to resources that are contained in g.

Similarly, relations between resources that are contained in the same group, and a *part-of-relation* between resources and groups can be deduced. In the following section we present ranking algorithms, which exploit these new relations.

4. RANKING STRATEGIES

In this section we present GroupMe! ranking strategies. All strategies are based on the FolkRank algorithm [7] and differ in the way GroupMe! tag assignments (which form a *4-uniform hypergraph*, cf. Definition 3), are exploited in the graph construction process. Figure 3.i shows the hypergraph formed by the tag assignments of the scenario in Figure 2 – we visualize such hypergraphs similarly to [4]. The challenge of adapting the FolkRank algorithm to GroupMe! folksonomies is to identify semantically appropriate strategies for constructing a *graph*, whose adjacency matrix serves as input for the PageRank-based FolkRank algorithm.

4.1 FolkRank Algorithm

The core idea of the FolkRank algorithm is to transform the hypergraph formed by the traditional tag assignments (see Definition 1) into an undirected, weighted tripartite graph $\mathbb{G}_{\mathbb{F}} = (V_{\mathbb{F}}, E_{\mathbb{F}})$, which serves as input for an adaption

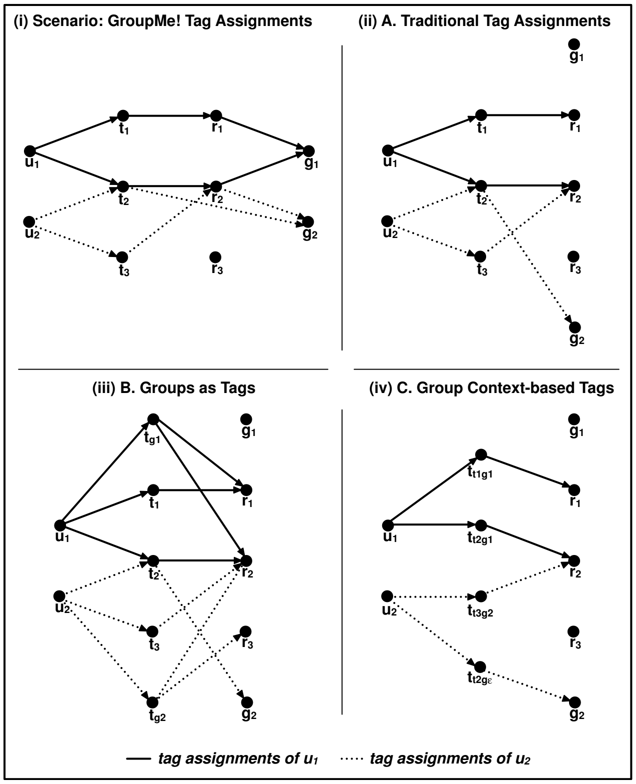

Figure 3: Strategies for interpreting the GroupMe! tag assignments from the scenario illustrated in Figure 2. The edges of the hypergraph (i) model GroupMe! tag assignments $(u, t, r, g) \in \breve{Y}$ whereas the edges of the hypergraphs (ii)-(iv) form triples (u, t, r) where $u \in U$, $t \in T$, and $r \in \breve{R}$.

of PageRank [14]. At this, the set of nodes is $V_{\mathbb{F}} = U \cup T \cup R$ and the set of edges is given via $E_{\mathbb{F}} = \{\{u, t\}, \{t, r\}, \{u, r\} | (u, t, r) \in Y\}\}$ (cf. Definition 1). The weight w of each edge is determined according to its frequency within the set of tag assignments, i.e. $w(u, t) = |\{r \in R : (u, t, r) \in Y\}|$ is the number of resources the user u tagged with keyword t. Accordingly, $w(t, r)$ counts the number of users who annotated resource r with tag t, and $w(u, r)$ determines the number of tags a user u assigned to a resource r. With $\mathbb{G}_{\mathbb{F}}$ represented by the real matrix A, which is obtained from the adjacency matrix by normalizing each row to have 1-norm equal to 1, and starting with any vector $\vec{w}$ of non-negative reals, PageRank iterates as follows:

$$\vec{w} \leftarrow dA\vec{w} + (1 - d)\vec{p}. \qquad (1)$$

PageRank utilizes vector $\vec{p}$ as a preference vector, fulfilling the condition $||\vec{w}||_1 = ||\vec{p}||_1$. Its influence can be adjusted by $d \in [0, 1]$. Based on this, FolkRank is defined as follows [7]:

Definition 4. The *FolkRank algorithm* computes a topic-specific ranking in folksonomies by executing the following steps:

1. $\vec{p}$ specifies the preference in a topic (e.g. preference for a given tag).
2. $\vec{w}_0$ is the result of applying the adapted PageRank with $d = 1$.
3. $\vec{w}_1$ is the result of applying the adapted PageRank with some $d < 1$.
4. $\vec{w} = \vec{w}_1 - \vec{w}_0$ is the final weight vector. $\vec{w}[x]$ denotes the *FolkRank* of $x \in V$.

4.2 FolkRank Adaptions

In order to adapt the FolkRank algorithm to GroupMe! folksonomies we confine ourself on adapting the process of constructing the graph $\mathbb{G}_{\mathbb{F}}$ from the hypergraph formed by the GroupMe! tag assignments. Therefore, we introduce three main strategies:

A. Traditional Tag Assignments This approach reduces GroupMe! tag assignments to traditional tag assignments, and therewith constructs a 3-uniform hypergraph as visualized in Figure 3.ii. Groups are just taken into account as resources which might or might not be tagged. Building the tripartite graph $\mathbb{G}_{\mathbb{A}} = (V_{\mathbb{A}}, E_{\mathbb{A}})$ is done analogously to FolkRank. The set of nodes and edges is given as follows: $V_{\mathbb{A}} = U \cup T \cup \breve{R}$ and $E_{\mathbb{A}} = \{\{u, t\}, \{t, r\}, \{u, r\} | u \in U, t \in T, r \in \breve{R}, g \in G \cup \{\varepsilon\}, (u, t, r, g) \in \breve{Y}\}$. Computing the weight of each edge also corresponds to the FolkRank approach, e.g.: $w(u, t) = |\{r \in \breve{R} : (u, t, r, g) \in \breve{Y}\}|$ is the number of resources the user u tagged with keyword t in any group g. This strategy corresponds to the *normal FolkRank* algorithm. It just requires the preprocessing step, in which the GroupMe! folksonomy is transformed into a traditional folksonomy.

B. Groups as Tags In this strategy we create artificial tags $t_g \in T_G$, $T_G \cap T = \emptyset$, for each group g and assign such tags to all resources contained in g, whereby the user who added a resource to the group, is declared as the *tagger*. The set of nodes is extended by T_G: $V_{\mathbb{B}} = V_{\mathbb{A}} \cup T_G$. The edges added to $V_{\mathbb{F}}$ by the strategy are: $E_{\mathbb{B}} = E_{\mathbb{A}} \cup \{\{u, t_g\}, \{t_g, r\}, \{u, r\} | u \in U, t_g \in T_G, r \in \breve{R}, u \text{ has added r to group } g\}$. We use a constant value w_c to weight these edges because a resource is usually added only once to a certain group. Hence, counting e.g. the number of users who added a resource to a specific group would not make sense[6]. The hypergraph, which functions as database for this graph construction strategy, is depicted in Figure 3.iii. Here, for both groups g_1 and g_2, which are treated as normal resources, two new artificial tags t_{g_1} and t_{g_2} are introduced and assigned to those resources which are member of the corresponding group.

C. Group Context-based Tags If users assign a certain tag to resources in different groups then the meaning of the tag may differ. This strategy replaces every tag t with a tag t_{tg}, which indicates that tag t was used in group g. It then transforms all GroupMe! tag assignments into normal tag assignment triples. For example, the GroupMe! tag assignment (u_1, t_2, r_2, g_1), presented in Figure 2, is interpreted as $(u_1, t_{t_2 g_1}, r_2)$ ($= tas_1$). Assume we also have a tag assignment (u_1, t_2, r_2, g_2) then this would be converted into $(u_2, t_{t_2 g_2}, r_2)$ ($= tas_2$). Thus, a 3-uniform hypergraph is build, which

[6]Instead we select e.g. $w_c(t_g, r) \approx max(|\{u \in U : (u, t, r, g) \in \breve{Y}\}|)$ as we believe that grouping a resource is in general more valuable than tagging it.

serves as input for the construction of $\mathbb{G}_{\mathbb{C}}$. The construction of $\mathbb{G}_{\mathbb{C}}$ is done as in the normal FolkRank algorithm, described in Section 4.1. Detecting equality of tags is the only exception, e.g. given tas_1 and tas_2 from above, the weight $w(u_1, t_{t_2 g_1})$ is not only determined by tas_1 but also partially by tas_2, although the tag $t_{t_2 g_1}$ in tas_1 is not exactly equal to $t_{t_2 g_2}$ in tas_2. We compute the similarity between two tags $t_{t_x g_y}$ and $t_{t_v g_w}$ and therewith the influence of a tag assignment to a weight as follows:

$\wedge$	$t_x = t_v$	$t_x \neq t_v$
$g_y = g_w$	1.0	0.2
$g_y \neq g_w$	0.4	0

Hence, based on tas_1 and tas_2 it is $w(u_1, t_{t_2 g_1}) = 1.4$.

In addition to the three strategies that can be applied to generate the graph $\mathbb{G}$, which serves as input for the FolkRank algorithm, we present two further strategies to exploit a GroupMe! folksonomy. They can be applied as extensions to the strategies above. The core idea of both strategies is to propagate tags assigned to one resource/group to other resources or groups. Such techniques synthetically increase the amount of input data and do not require to change the strategies described above substantially.

Propagation of Group Tags GroupMe! users tag groups about 1.75 times more often than common resources [2]. By propagating tags which have been assigned to a group (*group tags*) to its resources we try to counteract this situation. For example in Figure 2, tag t_2, which is assigned to group g_2, can be propagated to all resources contained in g_2. An obvious benefit of this procedure is that untagged resources like r_3 obtain tag assignments (here: (u_2, t_2, r_3, g_2)). In order to adjust the influence of inherited tag assignments, we weight these assignments by a dampen factor $df \in [0, 1]$.

Propagation of all Tags In the same way tags can be propagated among resources that are contained in the same group. This strategy induces propagation of (i) group tags to resources within the group, (ii) resource tags of one resource to other resources within a group, and (iii) resource tags to the group itself. Note that only tag assignments that have been carried out within the context of the corresponding group are considered for propagation.

5. EVALUATIONS

In this section we combine the strategies A-C with the proposed options for propagating tags and evaluate which of the combinations works best.

5.1 Test Setting

For the evaluations, presented in this section, we used a snapshot of the GroupMe! data set, which consists of 235 users, 978 tags, 1351 resources, 273 groups, and 1758 tag assignments. The groups contain in average 4.79 resources, and groups were tagged 1.75 times more often than resources. Furthermore, 565 resources and 101 groups have not been tagged with any keyword.

According to the *tagging system design taxonomy* proposed in [10], GroupMe! is a free-for-all tagging system, which allows users to annotate multimedia content for future retrieval. Hence, GroupMe! allows for broad folksonomies as every user is allowed to tag every resource or group without any restrictions. Tagging a resource $r \in \check{R}$ is done when users are situated in the view of a certain group g. Thereby, users are only able to see those tags that have been assigned to r within the context of the group g (same holds for group g). Explicit tag suggestions are not provided by the GroupMe! system. However, the tag cloud of a group and the resource's visualization, which is adapted to the media type of the resource, help the users to reflect on appropriate tags for the resource. For a more detailed description of the GroupMe! tagging design the reader is referred to [2].

The adapted FolkRank algorithms compute rankings for all entities of a folksonomy (users, tags, and resources). In the evaluation we concentrate on ranking of resources and groups as search for resources is the most common use case of ranking in folksonomy systems. In order to measure the quality of our ranking strategies we used the $OSim$ and $KSim$ metrics as proposed in [5]. $OSim(\tau_1, \tau_2)$ enables us to determine the overlap between the top k resources of two rankings, τ_1 and τ_2.

$$OSim(\tau_1, \tau_2) = \frac{|R_1 \cap R_2|}{k}, \tag{2}$$

where $R_1, R_2 \subset \check{R}$ are the sets of resources that are contained in the top k of ranking τ_1 and τ_2 respectively, and $|R_1| = |R_2| = k$.

$KSim(\tau_1, \tau_2)$, which is based on Kendall's τ distance measure, indicates the degree of pairwise distinct resources, r_u and r_v, within the top k that have the same relative order in both rankings.

$$KSim(\tau_1, \tau_2) = \frac{|\{(u, v) : \tau_1, \tau_2 \text{ agree on order of } (u, v), u \neq v\}|}{|U| * (|U| - 1)} \tag{3}$$

U is the union of resources of both top k rankings. τ_1' corresponds to ranking τ_1 extended with resources R_1' that are contained in the top k of τ_2 and not contained in τ_1. We do not make any statements about the order of resources $r \in R_1'$ within ranking τ_1'.t τ_2' is constructed correspondingly.

Together, *OSim* and *KSim* are suited to measure the quality of a ranking with respect to an optimal (possibly hand-selected) ranking. Our evaluations are based on 50 hand-selected rankings: Given 10 keywords, which were out of T, and the entire GroupMe! data set, 5 experts independently created rankings for each of the keywords, which represented from their perspective the most precise top 20 ranking. By building the average ranking for each keyword, we gained 10 optimal rankings. Among the 10 keywords, there are frequently used tags as well as seldom used ones.

5.2 Measurements and Discussion

Table 5.2 gives an overview on the measured results for each ranking strategy introduced in Section 4 with respect to OSim and KSim metrics. The strategies are ordered according to their OSim performance, whereas both, OSim and KSim values are averaged out of 10 test series (for the 10 different keywords and corresponding hand-selected rankings). In terms of the OSim, "C – Group Context-based Tags" can be identified as best strategy: It computes rankings,

Rank	Hand-selected	A. Traditional TAS	B. Groups as Tags	C + Full Tag Propagation
1.	Optimizing web search using social annotations	Optimizing web search using social annotations	Optimizing web search using social annotations	Yahoo! research
2.	Exploring social anno tations for the sem...	The Semantic Web: Will It All End In Tiers?	HITS	Optimizing web search using social annotations
3.	Personalized PageRank	New *Semantic* Web!	Webpage Ranking *(group)*	Ontologies are us
4.	SimRank	The Semantic Web: An Introduction	SimRank	Bibsonomy
5.	PageRank	The Semantic Web: Scientific American	PageRank	HITS
6.	FolkRank	LEGOLAND	Topic-sensitive PageRank	SimRank
7.	Ontologies are us	eschbach *(group)*	Personalized PageRank	PageRank
8.	Topic-sensitive PageRank	Andreas Eschbach - Wikipedia	Yahoo! research	Topic-sensitive PageRank
9.	Bibsonomy	Andreas Eschbach Homepage	FRank	Personalized PageRank
10.	FRank	Andreas Eschbach: Der Nobelpreis	Ontologies are us	FolkRank

Table 2: Top 10 rankings computed by different ranking strategies (and by hand respectively) for the term *"socialpagerank"*.

	Strategy	OSim	KSim
(i)	C + Full Tag Propagation	0.610	0.369
(ii)	B + Group Tag Propagation	0.585	0.368
(iii)	B	0.580	0.375
(iv)	C + Group Tag Propagation	0.540	0.351
(v)	B + Full Tag Propagation	0.465	0.273
(vi)	A	0.405	0.255
(vii)	C	0.390	0.257
(viii)	A + Group Tag Propagation	0.360	0.237
(ix)	A + Full Tag Propagation	0.345	0.247

Table 1: Overview of OSim and KSim for different ranking strategies ordered by OSim, where the dampen factor for propagating tags is 0.2. A denotes the *"Traditional Tag Assignments"* strategy, B is the *"Groups as Tags" strategy*, and C is the *"Group Context-based Tags"*

which contain 61% of the resources that also occur in the corresponding hand-selected top 20 ranking lists. Group tag propagation does not influence the approach "B – Groups as Tags" strongly as strategies (ii) and (iii) have nearly the same OSim values. This can be explained with the functionality of "B": For each group, "B" introduces artificial tags and assigns those tags to the group's members. Considering the graph structure, this almost conforms to propagating the tags of a group to its members.

Strategy (vi) does not exploit the group structure as it reduces GroupMe! tag assignments to traditional tag assignments (see Section 4.1) and can therewith be interpreted as the traditional FolkRank algorithm. The extensions of FolkRank, (viii) and (ix), which rudimentary exploit the group structure, do not improve the overlapping similarity of 0.405 but rather degrade the quality of FolkRank. We assume that the approach of propagating tags without modeling the group dimension within the graph, which serves as input for the ranking algorithm, primarily increases the recall but has a negative effect on the precision.

Regarding the KSim, strategy (iii), which treats groups as tags, performs best, followed by strategies (i), (ii), and (iv). The quality of the strategies (i)-(iv) is, in view of KSim, more than 30% better than the quality of strategies (v)-(ix).

Figure 4 gives an idea about how the ranking strategies behave when varying the dampen factor for tag propagation. Naturally, the dampen factor does not effect strategies

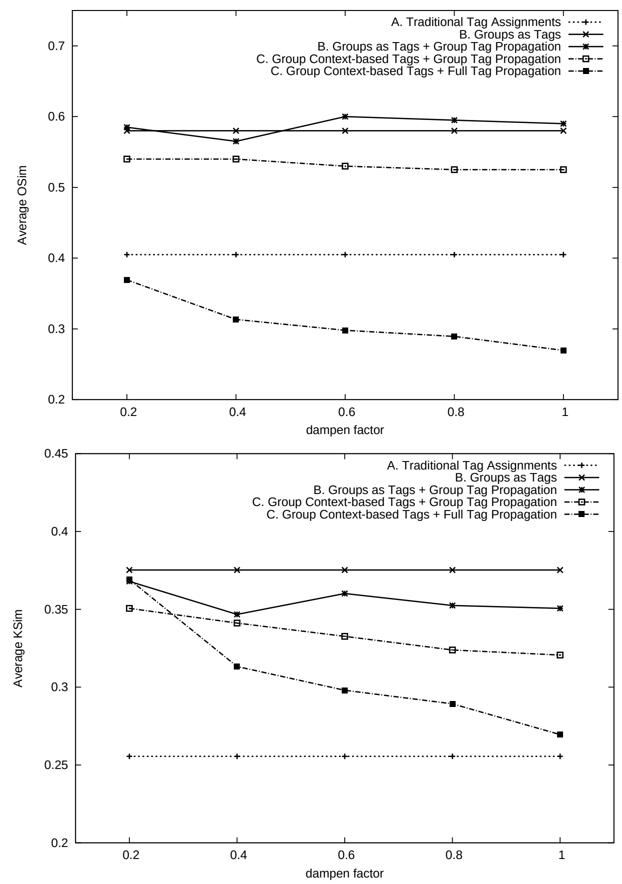

Figure 4: Average OSim and KSim (with respect to 10 different top 20 ranking comparisons) for varying dampen factors, which control the influence of propagated tags, and different ranking strategies.

"A – Traditional Tag Assignments" and "B – Groups as Tags" because both strategies do not make use of tag propagation. When varying the dampen factor, the OSim value is comparatively constant as well as for the strategies that base on propagation of group tags. The OSim and KSim of strategy "C + Full Tag Propagation" continuously degrades when the dampen factor increases. Gazing at the idea of "Full Tag Propagation" illustrates this behavior: Assume there is a resource r in a group g, which contains 20 other resources, and r is the only resource, which is tagged with

t. Then, propagation of t to g and all resources of g with a dampen factor of 1.0 would conceal the prominent role of resource r in terms of tag t. Hence, ranking the resources of g in an adequate order gets difficult (see KSim value), and the increased recall complicates the process of identifying resources to put into the top k of the ranking for tag t.

Table 5.2 outlines example rankings computed for the tag *"socialpagerank"* by different ranking strategies. Furthermore, it lists the corresponding hand-selected ranking, which is based on votings of five experts. Within the GroupMe! data set the resource *"Optimizing web search using social annotations"*, a paper which proposes the SocialPageRank algorithm, was the only resource tagged with the keyword "socialpagerank". This resource was ranked at first place in the hand-selected ranking, and almost every ranking strategy conforms to this decision. Starting from the second position the ranking of strategy "A", which represents the traditional FolkRank algorithm, gets imprecise. As strategy "A – Traditional Tag Assignments" does not exploit the group structure, the only solution to discover other relevant resources rests upon the users, who annotated the resource, and other tags that have been assigned to the resource. The group-based ranking strategies, on the other hand, are able to detect adequate resources via the group containing the resource.In the given example, this group is *"Webpage Ranking"* and strategy "B – Groups as Tags" is the only strategy that lists the group also within the top 10.

5.3 Results

The goal of our investigation was to identify whether grouping of resources in folksonomies has an impact on the quality of search strategies in social networks. To give proof on our hypothesis that grouping improves the quality of search, it is necessary to compare the search strategies which explore the grouping context to those search strategies which do not. As benchmark, we have chosen the FolkRank algorithm, and have developed search algorithms that extend FolkRank to exploit the group context as described in Section 4.1. All algorithms, FolkRank as well as the group-aware extensions, were tested under the same conditions, i.e. the same set of data, hardware, etc.

We tested our hypothesis with a one-tailed t-Test. The null hypothesis H_0 is that some group-aware FolkRank extension is as good as a the normal FolkRank without group-awareness, while H_1 states that some group-aware FolkRank extension is better than normal FolkRank. We tested it with a significance level of $\alpha = 0.05$. Tests were performed for the two measures OSim and KSim (see Section 5.1):

OSim With respect to OSim, the strategy "B – Groups as Tags" is significantly better than normal FolkRank.
Furthermore, FolkRank did not improve if we applied any of the tag-propagation strategies described in Section 4.2, and, indeed, the strategy "B – Groups as Tags" was significantly better than normal FolkRank with or without tag propagation. The variations of "B – Groups as Tags" to reflect tag propagation were, one compared to the other, not significantly different, but only the propagation of group tags showed significant improvement in comparison to FolkRank (with or without tag propagation).
Also the strategy "C – Group Context-based Tags", where full propagation of tags was performed (damping factor 0.2), was significantly better than the normal FolkRank regardless whether any propagation of tags was performed in the latter. From our test data, we hypothesize that strategy "C" benefits from the propagation of tags while "B" does not. Our actual data did not give statistically significant results on this, and we will investigate the impact on tag propagation in our future work.

KSim With respect to KSim, the strategy "B – Groups as Tags" is significantly better than normal FolkRank, whether or not the latter uses any tag propagation strategy.
Also the strategy "C – Group Context-based Tags", where group tags are propagated (damping factor 0.2) is significantly better than normal FolkRank, whether or not the latter uses any tag propagation strategy.

OSim *and* KSim Only the strategy "B – Groups as Tags" (without tag propagation or with group tag propagation, damping factor 0.2) was significantly better with respect to both measures, OSim and KSim, than normal FolkRank (whether or not the latter uses any tag propagation strategy).

Our evaluation indicated that the grouping of resources significantly improves the quality of search in folksonomies. The grouping activity itself brings many advantages for users: they can organize resources of interest, they can overlook and inspect a group's content, they can share groups with fellow users, and can explore the information in a folksonomy in novel ways, e.g. by requesting new, artificial groups that collect contents of all groups for the same topic, that collect the most popular groups or resource, etc. Furthermore, the drag & drop metaphor realized in the GroupMe! system makes the grouping activity intuitive for users, and from our experience with running GroupMe! we have seen that users like grouping [2]. Thus, while grouping is an easy and well-received feature for folksonomies, this activity provides, on the technical side, valuable information to detect relevant resources, and to improve the quality of search, and seems to be a very promising approach to improve current folksonomy systems.

6. RELATED WORK

In this paper, our main motivation was not to find the best possible search algorithm for folksonomies or to improve already existing search algorithms. Instead, our motivational question to be answered is if additional context information, relating resources with each other, namely the group concept, which is special for the GroupMe! system [1], can be exploited to improve the search performance in folksonomies. For this comparison we found that the FolkRank[7] algorithm, which is used in the BibSonomy system [6], is a perfect base to be extended: The algorithm mainly consists of two parts: First, the algorithm reduces the hypergraph to a two dimensional graph which afterwards is processed by a PageRank variant, called the adapted PageRank. This enables us to adapt to the FolkRank in different extends: a) adapting the construction of the input graph from the hypergraph spanned by the tag assignments or b) modifying the adapted PageRank and FolkRank algorithm respectively. As any adaption of the part b) would make it hard to prove that the performance gain was resulted by the new group structure instead of the tuning of the FolkRank algorithm itself, we decided to adapt only the process of transforming the hypergraph into a graph. Hence, we can study on the group effect clearly. Furthermore, our different strategies can be easily used to construct the input graph for other folksonomy-based search algorithms, like SocialPageRank or SocialSimRank [3], where the latter is an adaption of the SimRank algorithm [8].

We tested our ranking strategies with a snapshot of the GroupMe! data set because there are, to the best of our knowledge, no other tagging systems that gain data comparable to the GroupMe! folksonomy. Flickr enables users to create groups of images but does not allow to tag those groups. Social bookmarking systems like BibSonomy or Connotea[7] provide functionality to create groups of users, however they do not offer functionality to structure bookmarks into groups. Creating a GroupMe!-like data set synthetically, e.g. based on del.icio.us bundles, is not appropriate because grouping as well as tagging is an activity done by users.

The design of a tagging systems has an important impact on resulting folksonomies. In [10] Marlow et al. propose a formal model (*tagging design taxonomy*) to classify tagging systems. According to the tagging design taxonomy, GroupMe! and BibSonomy are similar in many aspects (e.g. regarding the *tagging rights*: every user is allowed to tag everything), which was another motivation to adapt the FolkRank algorithm. The impact of the tagging system design on resulting folksonomies as well as the performance of ranking strategies with respect to different kind of folksonomies – e.g. broad vs. narrow folksonomies [11] – are open issues to analyze.

7. CONCLUSIONS

Folksonomies are characterized by a bottom-up approach to knowledge creation: Many people leave their traces, annotate resources, share annotations with others, browse via annotations, annotate again, etc. They are an interesting and highly dynamic source of information, and bear great potential for information discovery and retrieval. With the GroupMe! system we have created an intuitive Web 2.0 system that allows users to organize and maintain Web resources very easily. The system offers - like other current Web 2.0 systems - the tagging feature, but in addition enables users to group Web resources they consider interesting together, and tag also the groups. We capture the semantics of user interactions with the GroupMe! system, and exploit the dynamic and evolving grouping information for search. To verify that the grouping information improves the quality of ranking in folksonomy systems, we compared the results of ranking algorithms which can reflect the group information to those which cannot. To realize such a comparison, we benchmarked our algorithms against the FolkRank algorithm [7], and are able to show that grouping of resources in folksonomies significantly improves the quality of search result ranking.

8. ACKNOWLEDGMENTS

We thank Mischa Frank and Patrick Siehndel for their contribution and engagement in realizing the GroupMe! system.

[7]http://connotea.org

9. REFERENCES

[1] F. Abel, M. Frank, N. Henze, D. Krause, D. Plappert, and P. Siehndel. GroupMe! – Where Semantic Web meets Web 2.0. In *Int. Semantic Web Conference (ISWC 2007)*, November 2007.

[2] F. Abel, N. Henze, and D. Krause. A Novel Approach to Social Tagging: GroupMe! In *4th Int. Conference on Web Information Systems and Technologies (WEBIST)*, May 2008.

[3] S. Bao, G. Xue, X. Wu, Y. Yu, B. Fei, and Z. Su. Optimizing Web Search using Social Annotations. In *WWW '07: Proceedings of the 16th Int. World Wide Web Conference*, pages 501–510, New York, USA, 2007. ACM Press.

[4] H. Halpin, V. Robu, and H. Shepherd. The Complex Dynamics of Collaborative Tagging. In *WWW '07: Proc. of the 16th Int. World Wide Web Conference*, pages 211–220, New York, USA, 2007. ACM Press.

[5] T. H. Haveliwala. Topic-Sensitive PageRank: A Context-Sensitive Ranking Algorithm for Web Search. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):784–796, 2003.

[6] A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme. BibSonomy: A Social Bookmark and Publication Sharing System. In A. de Moor, S. Polovina, and H. Delugach, editors, *Proc. First Conceptual Structures Tool Interoperability Workshop*, pages 87–102, Aalborg, 2006.

[7] A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme. FolkRank: A Ranking Algorithm for Folksonomies. In *Proc. FGIR 2006*, 2006.

[8] G. Jeh and J. Widom. SimRank: A Measure of Structural-Context Similarity. In *Proc. of Int. Conference on Knowledge Discovery and Data Mining (SIGKDD)*, Edmonton, Alberta, Canada, July 2002. ACM Press.

[9] B. Lund, T. Hammond, M. Flack, and T. Hannay. Social Bookmarking Tools (II): A Case Study - Connotea. *D-Lib Magazine*, 11(4), April 2005.

[10] C. Marlow, M. Naaman, D. Boyd, and M. Davis. HT06, tagging paper, taxonomy, flickr, academic article, to read. In *HYPERTEXT '06: Proc. of the Seventeenth Conference on Hypertext and Hypermedia*, pages 31–40, New York, USA, 2006. ACM Press.

[11] T. Vander Wal. Explaining and showing broad and narrow folksonomies. *http://www.personalinfocloud.com/2005/02/explaining_and_.html*, February 2005.

[12] T. Vander Wal. Folksonomy. *http://vanderwal.net/folksonomy.html*, July 2007.

[13] P. Mika. Ontologies Are Us: A unified model of social networks and semantics. In *Proc. International Semantic Web Conference*, pages 522–536, Nov. 2005.

[14] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank Citation Ranking: Bringing Order to the Web. Technical report, Stanford Digital Library Technologies Project, 1998.

[15] X. Wu, L. Zhang, and Y. Yu. Exploring Social Annotations for the Semantic Web. In *WWW '06: Proc. of the 15th Int. World Wide Web Conference*, pages 417–426, New York, NY, USA, 2006. ACM Press.