A NOVEL APPROACH TO SOCIAL TAGGING: GROUPME!
Enhancing Social Tagging Systems with Groups

Fabian Abel, Nicola Henze and Daniel Krause

IVS – Semantic Web Group, Leibniz University Hannover, Appelstr. 4, 30167 Hannover, Germany

Keywords: Social Media, Semantic Web, Tagging, Folksonomy, GroupMe!.

Abstract: Common social tagging systems like Flickr, del.icio.us and others lately became very popular. The key benefits of these systems include that users get involved in the content creation process, can easily — without overhead – comment or annotate Web content, share resources with fellow users, and benefit from the comments / annotations of other users with improved retrieval support. With GroupMe! we extend the idea of current social tagging systems by enabling users to not only tag Web resources they are interested in, but also to create collections (*groups*) of these Web resources by simple drag & drop operations. The grouping metaphor is intuitive and easy for the users, and our evaluation shows that users appreciate the grouping facility, and use this feature to organize and structure Web content. Technically, the grouping of resources carries valuable information about Web resources and their relations, and can be exploited to improve the mining of Web content, e.g. for search and retrieval.

1 INTRODUCTION

Popular systems like Flickr[1], YouTube[2], Blogger[3] or others, which allow users to share photos, broadcast own videos, or blog about topics they are interested in, are obvious indicators for the success of Web 2.0. These systems have shown that Web users are not absolutely satisfied with their role of pure *content consumers*. Instead, Web users want to contribute and collaborate actively by providing their own content, or by annotating content of other users. Web 2.0 websites fulfill these needs perfectly: While the AJAX technique (Garrett, 2005) enables websites to be interactive, lots of Web 2.0 websites allow users to create, upload or annotate content with free chosen keywords (so-called *tags*) collaboratively, hence enable *intercreativity*[4]. *Tagging* is one of the key factors that makes these systems successful: It allows users to annotate content with relevant keywords for future retrieval. Especially multimedia content becomes in this way searchable. With tagging, users create metadata *collaboratively*.

Another very well received feature of Web 2.0 systems is the sharing of favorite resources with fellow users. E.g., in YouTube, users can store their favorite videos, in BibSonomy[5] their favorite academic papers, and in del.icio.us[6] their favorite bookmarks. However, all of these systems are more or less limited to a certain media type: Some systems (like YouTube) support only one media type (videos), while other systems, which can handle different kind of media types, lack of an appropriate visualization (in del.icio.us, e.g., all media types are displayed as normal text links).

In this paper, we present the GroupMe! system[7] which offers a novel user interface to organize multimedia Web resources. The core idea of the GroupMe! approach is that users can group – via drag & drop – the Web resources they are interested in. Appropriate media wrappers ensure that content of groups is displayed in a concise manner. We report about the evaluation of the systems which shows that a) the GroupMe! group concept is very well accepted by the users, b) that users like to combine resources of different media types, and c) that these groups can be used to improve search.

The paper is structured as follows: In Section 2

---

[1]http://www.flickr.com/
[2]http://www.youtube.com/
[3]http://www.blogger.com/
[4]http://www.w3.org/Talks/9602seybold/slide6.htm

[5]http://www.bibsonomy.org/
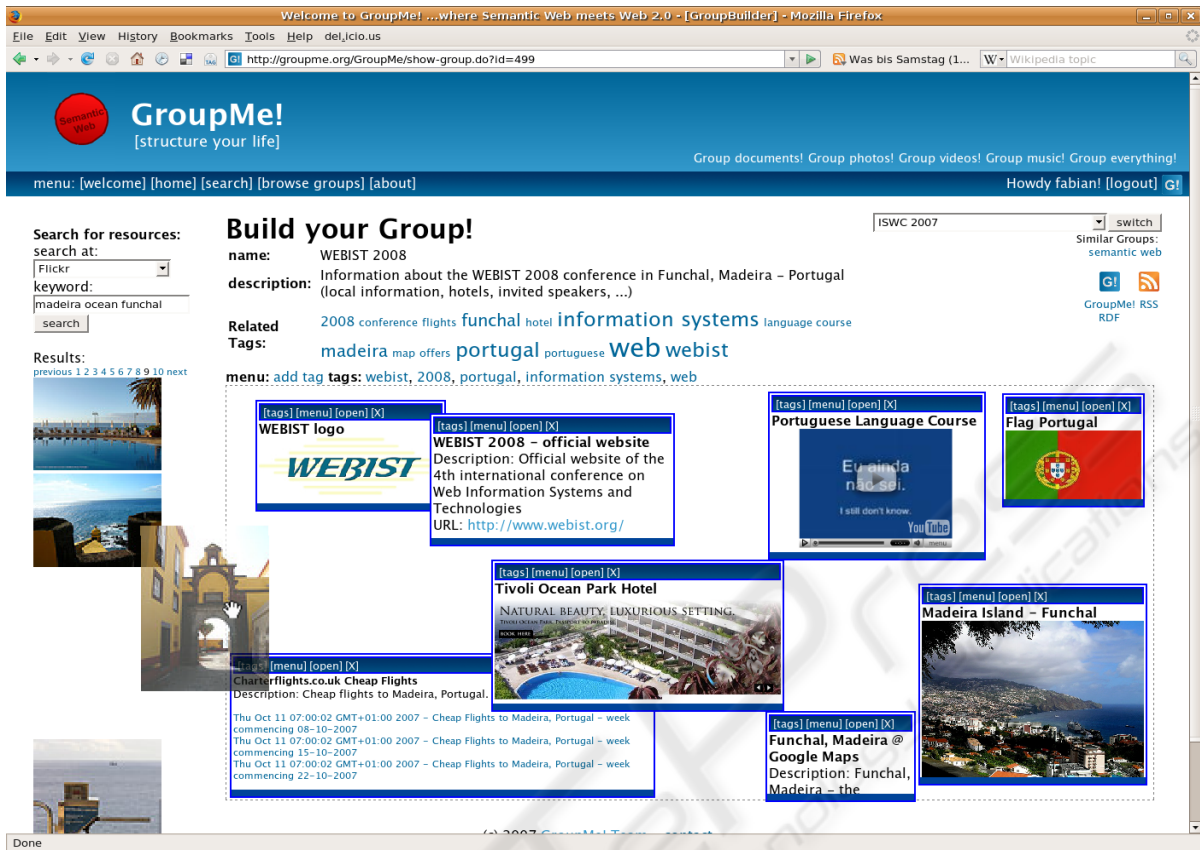[6]http://del.icio.us/
[7]http://groupme.org/

Figure 1: Screenshot of GroupMe! application: A user drags a photo from the left-hand side Flickr search bar into the GroupMe! group on the right-hand side.

we describe our GroupMe! system which is evaluated in Section 3. In Section 4 we compare the GroupMe! system with other state-of-the-art tagging systems and end with the Conclusion in Section 5.

2 GROUPME! SYSTEM

GroupMe! is a new kind of resource sharing system. It is comparable to social bookmarking systems – like del.icio.us – as it enables users to bookmark Web resources and annotate them with free chosen keywords. The core idea of GroupMe! is that users build groups of arbitrary multimedia Web resources on a specific topic and tag both, Web resources and groups. Groups can be understood as lightweight wiki pages. But instead of writing own content, users create groups via simple *drag & drop* operations and via visual arrangement of contained resources.

Figure 1 shows a screenshot of the GroupMe! system. It illustrates a typical scenario. Let us assume that user *fabian* plans a trip to the *WEBIST 2008* conference in Funchal, Portugal. Therefore, he wants to build a GroupMe! group containing resources that are

relevant for the trip. Building such a group is simple and requires just three steps. At first fabian specifies the group's name (WEBIST 2008), then he utilizes integrated search engines – like Google or Flickr search engine – in order to search for adequate resources, and finally he adds resources, which are from his point of view relevant, via drag & drop into the group. Figure 1 depicts such a drag & drop operation with a photo gathered from Flickr. Furthermore, it shows the entire group fabian has designed. This group contains images (like a photo of Madeira Island), the official website of the WEBIST 2008, and a shockwave flash movie, which presents several photos of *Tivoli Ocean Park Hotel* – where the conference will be held – and gives the opportunity to book rooms. All elements are visualized according to their media types so that fabian and other users can see relevant information at a glance. For example the RSS news feed that informs about cheap flights to Madeira (see bottom left) directly presents the latest flight offers to the user. And the video showing a Portuguese language course (see top right) can be played back immediately.

Altogether the arranged group in Figure 2 appears

like a *collage of information artifacts* about *WEBIST 2008 trip*, which is comprehensible for users. Importantly, content of this group is also accessible and understandable for machines. Because, when users create groups, GroupMe! produces RDF. This is done in two different ways:

1. Each user interaction – grouping and tagging – is captured as RDF using several vocabularies, e.g. *FOAF*[8] and a GroupMe!-specific vocabulary[9] that defines new GroupMe! concepts.

2. Whenever a user drops a new Web resource into a group, domain dependent content extractors gather useful metadata so that resources can be enriched with semantically well defined descriptions. When e.g. adding a Flickr photo into a group, GroupMe! translates Flickr-specific descriptions into a well defined RDF description using *DCMI element set*[10].

RDF created in GroupMe! is made available to other Web applications and can be accessed via RSS feeds or RESTful API. Hence, other applications can benefit from the feature of grouping and enriching resources with machine understandable semantics.

Regarding the above scenario, GroupMe! can be utilized by a Web service that searches for photos according to a location specified via geographic coordinates. With the semantic descriptions captured in groups like the one about *WEBIST 2008 trip*, which contains photos of Funchal as well as a link to Google maps (see bottom in Figure 1) equipped with the coordinates of Funchal, such a service is now able to retrieve photos by locations even if these photos are not directly annotated with geographic coordinates.

2.1 GROUPME! Architecture

In technical terms, GroupMe! is a modular Web application that adheres to the Model-View-Controller pattern. It is implemented using the J2EE application framework *Spring*[11]. Figure 2 illustrates the underlying architecture, which consists of four basic layers:

Aggregation. The aggregation layer provides functionality to search for resources a user wants to add into GroupMe! groups. Currently, GroupMe! supports Google, Flickr, and of course a GroupMe!-internal search, as well as adding resources by specifying their URL manually. *Content Extractors* allow us to process gathered resources in order to extract useful data and meta-

---

[8]http://xmlns.com/foaf/spec/
[9]http://groupme.org/rdf/groupme.owl
[10]http://dublincore.org/documents/dces/
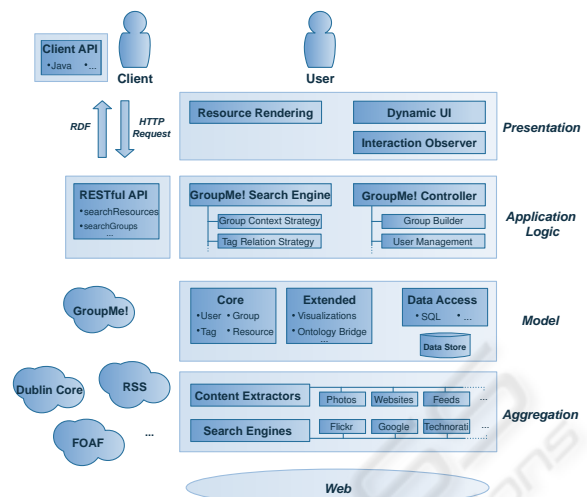[11]http://springframework.org/

Figure 2: Technical overview of the GroupMe! application.

data, which are converted to RDF using well-known vocabularies. As mentioned in Section 2, when e.g. adding a Flickr image into a group, a *Photo content extractor* converts Flickr-specific descriptions into RDF descriptions using Dublin Core vocabulary. At the moment of writing this paper, we are extending content extraction functionality by utilizing services like DBpedia[12] or Sindice[13], and frameworks like Aperture[14].

Model. The core GroupMe! model is composed of four main concepts: *User*, *Tag*, *Group*, and *Resource*. These concepts constitute the base for the GroupMe! folksonomy (cf. section 2.2). In addition, the model covers concepts concerning the users' arrangements of groups, etc. The *Data Access* layer cares about storing model objects. The actual data store backend is arbitrarily exchangeable. At the moment we are using a MySQL database.

Application Logic. The logic layer provides various controllers for modifying the model, exporting RDF, etc. The internal GroupMe! search functionality, which is implemented according to the strategy pattern in order to switch between different search and ranking strategies, is made available via a RESTful API. It enables third parties to benefit from the improved search capabilities (cf. Section 3.1), and to retrieve RDF descriptions about resources – even such resources that were not equipped with RDF descriptions before they were integrated into GroupMe!. To simplify

---

[12]http://dbpedia.org/
[13]http://sindice.com/
[14]http://aperture.sourceforge.net/

usage of exported RDF data, we further provide a lightweight Java *Client API*, which transforms RDF into GroupMe! model objects.

Presentation. The GUI of the GroupMe! application is based on AJAX principles. Therefore, we applied Ajax and JavaScript frameworks like script.aculo.us[15], DWR[16], or Prototype[17]. Such frameworks provide already functionality to drag & drop elements, resize elements, etc. Visualization of groups and resources is highly modular and extensible. Switching between components that render a specific resource or type of resource can be done dynamically, e.g. visualization of group elements is adapted to their media type (see Fig. 1). In the future, users should also be enabled to choose an appropriate resource visualization from different applicable options.

When creating or modifying groups, each user interaction (e.g. moving and resizing resources) is monitored and immediately communicated to the responsible GroupMe! controller so that e.g. the actual size or position of a resource within a group is stored in the database.

2.2 GROUPME! Folksonomy

In social tagging systems data is created by users (the folks), who assign freely chosen tags to resources ($\rightarrow$ *tag assignment*). The evolving collection of such tag assignments is called *folksonomy*[18]. In general, a folksonomy is formally defined using finite sets of users, tags and resources, and a finite set of tag assignments, whereas a tag assignment constitutes a triple of a certain user, tag and resource (cf. (Mika, 2007)). With GroupMe! we introduce a new concept to social tagging systems, namely groups.

Definition 1 (Group). *A group is a set of resources.* A group is a resource as well. Hence, groups can contain groups, and groups can be tagged by users. With definition 1 we extend the formal definition of a folksonomy introduced in (Hotho et al., 2006a) as follows.

Definition 2 (GroupMe! Folksonomy). *A GroupMe! folksonomy is a 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.*

- $\breve{R} = R \cup G$ *is the union of the set of resources and the set of groups.*
- $\breve{Y}$ *defines a relation between these sets (tag assignment):* $\breve{Y} \subseteq U \times T \times \breve{R} \times G$.

Thus, tagging of resources within the GroupMe! system is always done in context of a group, which itself may have tags. In comparison to traditional folksonomies, in which relations between tags mainly rely on their co-occurrences (i.e. two tags are assigned to the same resource), we obtain new relations between tags:

1. A relation between tags assigned to different resources that are contained in the same group.

2. A relation between tags assigned to a group and tags assigned to resources that are contained in the group.

Similarly, we gain relations between resources that are contained in the same group, and a *part-of-relation* between resources and groups. These new relations can be exploited by search and ranking algorithms. For example, when searching for resources with a given tag, an algorithm could also rank resources that are directly tagged with the given *query string* but which are member of a group that is tagged with that query string.

A more detailed discussion of such an algorithm and the GroupMe! folksonomy model can be found in (Abel et al., 2007).

3 EVALUATION

This section gives an analysis of the GroupMe! system, in particular on usage and tagging characteristics, and evaluates the effects of the structure given by the groups to search and retrieve resources. The data underlying the analysis was collected during the first three month after the system's launch on July 14, 2007. During the observed period, GroupMe! had a total of 502 resources of which 428 were normal resources and 74 (14.74%) were groups. Altogether, 929 tag assignments were monitored, with 1.85 tags per resource in average. The overall evolution of resources and groups is given in Figure 3.

Interestingly, groups were tagged more extensively than ordinary resources: In average, 2.53 tags were assigned to groups, whereas only 1.73 tags were attached to other resources. Thus, groups were tagged 1.5 times more often than traditional resources. This effect was present over time, as depicted in Fig. 4. Furthermore, at the end of the observed period only
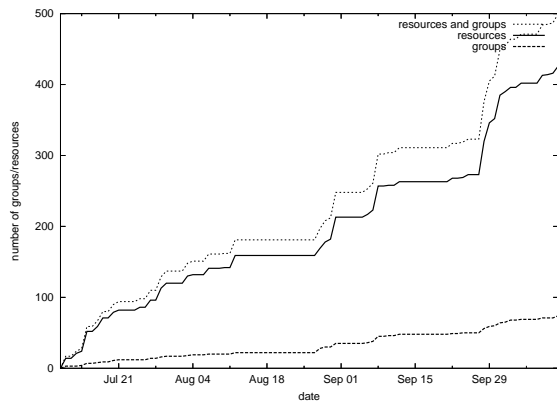
Figure 3: Evolution of number of resources/groups.

17.57% of the groups were not annotated with any tag in contrast to 32.71% of the resources. These initial observations give support for the hypothesis that users adopt the group idea to organize Web resources, and that they also invest in groups by annotating them.
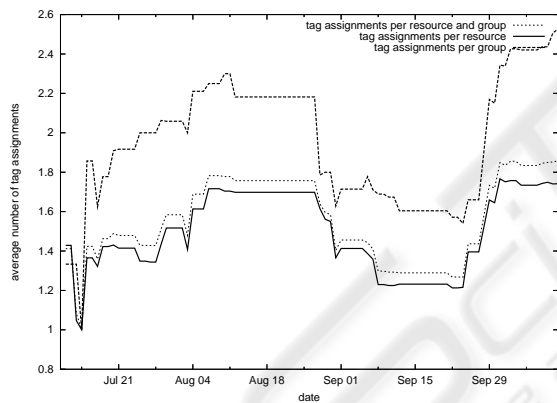

Figure 4: Average number of tags assigned to resources/groups.

A typical group in GroupMe! consists of 4 − 8 resources. That we do not observe groups with significantly more members can be explained from the user interface, which gives the users a canvas to place and arrange the Web resources. As the size of this canvas is limited, the on-screen display of the group becomes impractical with too many Web resources. Users collect resources with different media types in their group, as can be seen in table 1. Most popular among the media types are images, followed by videos and RSS feeds. Web sites, academic papers, presentation slides, etc. are denoted as *other Web resources* and are not mentioned separately, because to users they appear as simple bookmarks, i.e. their visualization is not yet adapted to their media type particularly. The possibility to include groups into a group

Table 1: Percentage of resources' media types that are part of GroupMe! groups.

Type of Resource	AVG Occurrences
images	41.01%
videos	8.57%
rss feeds	4.55%
groups	1.87%
other Web resources	43.96%

was only seldomly used, we explain this by the small number of available groups during the observation period.

3.1 Search Analysis

To examine the effect of groups on search in social tagging systems, we analyzed the search behavior of the users. A *search operation* in our experiment was either a search performed via the search interface, or a search which was initiated whenever a user clicked on a tag in the tag cloud. Obviously, not all clicks on a tag were intended to perform a search, but more often were used to explore the tags, and in particular the popular tags. Hence, we restricted our data set to only those search operations, which were followed by at least one click on a resource. Figure 5 shows how a search result looks like. Images are displayed directly, whereas groups, which are denoted by *"(Group)"*, and other resources are listed textual.

Figure 5: Screenshot: typical search result list.

We observed altogether 1747 search operations, to which in average 7.1 results were returned. 11.3% of the results delivered to a search were groups, and 53.11% of the returned results were not tagged with

the original search string or the tag on which the user had clicked (*query string*).

Table 2 lists the results of our experiments on search operations. We analyzed the top k search results of each search operation and users' first click on a particular resource within the top k. The average percentage of groups that were part of the top k results was between 10.38% and 15.39%. However, the percentage of search operations, in which users clicked on groups was between 19.40% and 24.51%. For example, when considering the top 10 search results, we observed that 12.15% of the results were groups but the percentage of *group clicks* was 1.84 times higher, namely 22.34%. Overall, normalization of resource and group clicks according to the number of resources and groups respectively reveals that groups were selected between 1.39 (Top 3) and 2.76% (Top 50) more frequent than ordinary resources. These observations support the hypothesis that groups itself constitute content users are interested in. The demand for groups is even higher than illustrated in Table 2 because groups can also be accessed using an explorative user interface different from the search interface. However, statistics of the explorative user interface, which has been utilized 271times, are not considered within our search analysis.

An important benefit of the GroupMe! system is that it provides the ability to increase the recall of queries. For our experiments we implemented search and ranking algorithms that take advantage of the GroupMe! folksonomy model in order to return also resources and groups that are not directly tagged with the given *query string* (cf. Section 2.2). In general, such *untagged resources* are ranked lower than resources which are directly tagged with the given query string. This explains the big increase of the percentage of untagged resources when increasing k (see *Untagged* in Table 2). Table 2 also points out that untagged resources are also well-accepted by users. Considering the top 10, 38.14% of the search results were not tagged with the query string and in 21.28% of the search operations users first clicked on an untagged resource, thus on a resource that would not have been found in a traditional tagging system which just considers direct resource annotations.

Consequently, the analysis of users' search operations validates two main hypotheses:

1. Users are interested in the new group concept: Groups are selected about two times more often than they occur within the search result list.

2. GroupMe!'s search and ranking strategies increase recall without reducing proportion of relevant resources remarkably: More than 20% of the clicks in the top 10 are performed on untagged re-

sources and groups.

4 RELATED WORK

GroupMe! is a social tagging system and competes with systems like BibSonomy, del.icio.us or Flickr. Table 3 summarizes some characteristics of GroupMe! according to the dimensions in the *tagging system design taxonomy* developed in (Marlow et al., 2006a), and compares them with related tagging systems.

Tagging Rights. GroupMe! allows every user to tag everything (*free-for-all*) as this enables us to gather more tags about a resource and alsoa higher variety of keywords than in constrainted systems. However, Flickr restricts tagging e.g. to the resource owner, friends, or contacts.

Tagging Support. When users annotate resources they are not supported with tag suggestions as this would limit the variety of tags. However, they have the ability to list tags that have already been assigned to a resource in context of the actual group. Tags, that have been assigned in context of other groups – and hence are possibly not appropriate in the actual group c007entext – are not visible to the user when tagging (*blind/viewable*).

Aggregation Model. In comparison to Flickr, which does not allow for duplicated tags (*set*), GroupMe! allows different users to assign the same tag to a certain resource (*bag*). This may enable a better evaluation of the importance of the tags.

Object Type. GroupMe! is the only system listed in Table 3 that supports tagging of resources displayed in a multimedia fashion. Although systems like del.icio.us enable users to bookmark and tag arbitrary Web resources, they just visualize resources in a textual way. Hence, while tagging e.g. an image in del.icio.us, users usually do not see the image they tag.

Source of Material. Resources that can be annotated and grouped in GroupMe! are globally distributed over the Web, and referenced by their URL. This enables GroupMe! to handle often changing resources like RSS feeds appropriately: Whenever a group is accessed, the most recent versions of the contained resources are displayed.

Social Connectivity. All systems listed in Table 3 allow users to be linked together. GroupMe! does not provide integrated features, but utilizes users' FOAF descriptions in order to identify links between users.

Table 2: Analysis of the Top k search results: (1) percentage of groups contained in the top k search results, (2) percentage of users who clicked on groups in the top k search results, (3) percentage of resources and groups in the top k that were not tagged with the given query string, (4) percentage of clicks on untagged resources and untagged groups.

Top k	Groups (1)	Group Clicks (2)	Untagged Resources (3)	Untagged Resource Clicks (4)
3	14.70%	19.40%	4.95%	14.93%
5	15.39%	21.69%	17.23%	15.66%
10	12.15%	22.34%	38.14%	21.28%
15	11.89%	22.45%	43.59%	23.47%
20	11.53%	23.00%	46.80%	25.00%
30	10.91%	24.51%	50.73%	25.49%
50	10.38%	24.27%	53.11%	26.21%

Table 3: GroupMe! tagging design in comparison to other social tagging systems. And user incentives in terms of tagging.

Dimension/System	GroupMe!	BibSonomy	del.icio.us	Flickr
Tagging Rights	free-for-all	free-for-all	free-for-all	permission-based
Tagging Support	blind/viewable	suggested	suggested	viewable
Aggregation Model	bag	bag	bag	set
Object Type	multimedia	textual	textual	images
Source of Material	global	global	global	user-contributed
Social Connectivity	links	links, groups	links	links
Resource Connectivity	groups	none	none	groups
5*User Incentives	- future retrieval - contribution - sharing - attract attention - self presentation	- future retrieval - contribution - sharing	- future retrieval - contribution - sharing - attract attention	- future retrieval - contribution - sharing - attract attention - self presentation

Resource Connectivity. Independent of the users' tags, a few resource sharing systems provide other features to connect resources. There are some systems that allow users to organize themselves into groups, and that provide functionality to retrieve resources, which are related to these groups – e.g. BibSonomy or Connotea[19]. However, to the best of our knowledge, Flickr and GroupMe! are at the moment the only notable tagging systems that enable users to assign resources to groups explicitly. Such hand-selected groups are highly valued by the users as indicated in our analysis (see Section 3.1).

User Incentives. GroupMe! users have several motivations to annotate resource ranging from simplification of future retrieval to self presentation (e.g. some users tag resources with *holiday* in order to express which locations they have visited).

What makes GroupMe! unique is that groups can be tagged and resources are always tagged in context of a specific group. Thereby, GroupMe! extends the traditional folksonomy model, which has been theorized in (Marlow et al., 2006b) or (Mika, 2007), and formalized in (Hotho et al., 2006a). With the GroupMe!

folksonomy model (see Section 2.2) new relations between resources, groups and tags emerge that can be exploited by search and ranking algorithms (Abel et al., 2007). Search and ranking algorithms that operate on traditional folksonomies have already been successfully applied in order to improve Web search. In (Bao et al., 2007) the authors introduced *Social-SimRank*, which adapts *SimRank* (Jeh and Widom, 2002) and computes similarity between tags and resources respectively. Furthermore, Bao et al. presented the *SocialPageRank* algorithm, which ranks Web resources according to how popular they are annotated. *FolkRank* (Hotho et al., 2006c) is another folksonomy-based search algorithm, which adapts the famous *PageRank* (Page et al., 1998) algorithm and involves user preferences. In our future work we will compare our algorithms, which exploit the GroupMe! folksonomy model, with the mentioned algorithms.

Learning relations between tags is another challenge in social tagging systems that can be utilized to improve retrieval of resources additionally. Hotho et al. presented an approach to mine association rules in folksonomies that point to subtag-supertag relations (Hotho et al., 2006b). The GroupMe! folksonomy model provides a foundation to deduce such relations more precisely, e.g. by analyzing tags that

---

[19]http://www.connotea.org/

have been assigned to a group and tags of group members. In (Rattenbury et al., 2007) the authors investigated how to learn more concrete semantics from folksonomies. In particular, they presented an approach to distinguish between *event tags* and *place tags*. Mentioned approaches for learning semantics can also be applied to GroupMe!. At the moment, instead of learning vague semantics, GroupMe! extracts semantic descriptions explicitly when new resources are added to a group. Hence, these descriptions can be utilized by machines offhand in order to search for certain type of resources (cf. example in Section 2). Therefore, all RDF produced in GroupMe! is fed back to the Web via RDF feeds and RESTful API. Other systems like *CiteULike*[20] or *BibSonomy* just offer RSS export.

5 CONCLUSIONS

GroupMe! gives users the possibility to *group* Web resources in an easy way – by simple drag & drop operations – and combines this idea with features of social tagging systems. The evaluation of GroupMe! shows that users appreciate the grouping facility to organize Web resources they are interested in. Groups can be seen as *hand selected* collections of Web content for a certain topic or domain. As such, they are also valuable results to search queries, and our investigations have shown that users recognize this and select groups among the search results often.

The structure inherently given by the groups can also be used to infer information about the content of Web resources. This is interesting for non-tagged Web resources, and particularly for multimedia Web resources whose content is - without tags - hardly determinable (like videos etc.). The analysis of the search behavior of users has revealed that this exploitation of grouping information uncovers relevant content, which – with tagging alone – would not have been found.

REFERENCES

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

Bao, S., Xue, G., Wu, X., Yu, Y., Fei, B., and Su, Z. (2007). Optimizing web search using social annotations. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 501–510, New York, NY, USA. ACM Press.

Garrett, J. J. (2005). AJAX: A new approach to web applications. *Adaptive Path. http://www.adaptive-path.com/publications/essays/archives/000385.php*.

Hotho, A., Jäschke, R., Schmitz, C., and Stumme, G. (2006a). BibSonomy: A social bookmark and publication sharing system. In de Moor, A., Polovina, S., and Delugach, H., editors, *Proc. First Conceptual Structures Tool Interoperability Workshop, 14th Int. Conf. on Conceptual Structures*, pages 87–102, Aalborg.

Hotho, A., Jäschke, R., Schmitz, C., and Stumme, G. (2006b). Emergent semantics in BibSonomy. In Hochberger, C. and Liskowsky, R., editors, *Informatik 2006 - Informatik für Menschen*, volume 94(2) of *LNI*, Bonn. GI.

Hotho, A., Jäschke, R., Schmitz, C., and Stumme, G. (2006c). FolkRank: A ranking algorithm for folksonomies. In *Proc. FGIR 2006*.

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

Marlow, C., Naaman, M., Boyd, D., and Davis, M. (2006a). HT06, tagging paper, taxonomy, flickr, academic article, to read. In *HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia*, pages 31–40, New York, NY, USA. ACM Press.

Marlow, C., Naaman, M., Boyd, D., and Davis, M. (2006b). Position Paper, Tagging, Taxonomy, Flickr, Article, ToRead. In *Collaborative Web Tagging Workshop at WWW2006*.

Mika, P. (2007). 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.

Page, L., Brin, S., Motwani, R., and Winograd, T. (1998). The Pagerank citation ranking: Bringing order to the web. Technical report, Stanford Digital Library Technologies Project.

Rattenbury, T., Good, N., and Naaman, M. (2007). Towards automatic extraction of event and place semantics from flickr tags. In *SIRIR '07: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 103–110, New York, NY, USA. ACM Press.

---

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